


MIT/WHOI 95-14

**Massachusetts Institute of Technology
Woods Hole Oceanographic Institution**



**Joint Program
in Oceanography/
Applied Ocean Science
and Engineering**



DOCTORAL DISSERTATION

An Entropic Framework for AUV Sensor Modelling

by

Hanumant Singh

June 1995


19960327 092

DTIC QUALITY INSPECTED 1

DISTRIBUTION STATEMENT A


Approved for public release;
Distribution Unlimited

DISCLAIMER NOTICE



THIS DOCUMENT IS BEST
QUALITY AVAILABLE. THE COPY
FURNISHED TO DTIC CONTAINED
A SIGNIFICANT NUMBER OF
PAGES WHICH DO NOT
REPRODUCE LEGIBLY.

DISCLAIMER NOTICE



THIS DOCUMENT IS BEST QUALITY AVAILABLE. THE COPY FURNISHED TO DTIC CONTAINED A SIGNIFICANT NUMBER OF COLOR PAGES WHICH DO NOT REPRODUCE LEGIBLY ON BLACK AND WHITE MICROFICHE.

MIT/WHOI

95-14

An Entropic Framework for
AUV Sensor Modelling

by

Hanumant Singh

Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

and

Woods Hole Oceanographic Institution
Woods Hole, Massachusetts 02543

June 1995

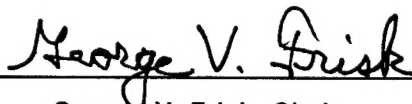
DOCTORAL DISSERTATION

Funding was provided by the Office of Naval Technology under
Grant No. N00014-90-C-0098.

Reproduction in whole or in part is permitted for any purpose of the United States
Government. This thesis should be cited as: Hanumant Singh, 1995. An Entropic
Framework for AUV Sensor Modelling. Ph.D. Thesis. MIT/WHOI, 95-14.

Approved for publication; distribution unlimited.

Approved for Distribution:



George V. Frisk, Chair

Department of Applied Ocean Physics and Engineering



Sallie W. Chisholm
MIT Director of Joint Program



John W. Farrington
WHOI Dean of Graduate Studies

An Entropic Framework for AUV Sensor Modelling

by

Hanumant Singh

B.S., Computer Science, Distinguished Graduate, George Mason University,
1989.

B.S., Electrical Engineering, Distinguished Graduate, George Mason University,
1989

Submitted to the Joint Committee on Oceanographic
Engineering in partial fulfillment of the requirements for
the degree of

Doctor of Philosophy

at the

Woods Hole Oceanographic Institution

and the

Massachusetts Institute of Technology

May 25, 1995

© 1995 Hanumant Singh. All Rights Reserved.

The author hereby grants to MIT and WHOI permission to reproduce and distribute publicly paper and electronic copies of this thesis document in whole or in part.

Author Hanumant Singh 5-25-'95
Joint Program in Oceanographic Engineering
Date Submitted

Certified by W.K. Stewart
W.K. Stewart, Thesis Supervisor
Assistant Scientist,
Applied Ocean Science and Engineering, WHOI

Accepted by Arthur B. Baggeroer
Professor A. Baggeroer
Chairman
Joint Committee on Oceanographic Engineering

An Entropic Framework for AUV Sensor Modelling

Hanumant Singh

Abstract

This thesis examines the general task of active sensing by defining a measure of efficiency for sensing in a particular environment. We focus on fine-scale acoustic mapping from an autonomous underwater vehicle (AUV). The constraints on imaging underwater — vehicle power, vehicle hydrodynamics, computational and telemetry requirements, and typical navigational and attitudinal uncertainties along with the underlying physics of the acoustic sensing modality — are considered in defining an entropic measure of sensor efficiency.

675-kHz pencil-beam sonar data acquired using the JASON remotely operated vehicle in a challenging shallow water environment and 200-kHz echo-sounder data acquired using the ABE AUV are used to demonstrate the utility of the entropic framework.

We show the utility of an entropic framework for the following: (i) Optimizing the speed of the AUV for maximizing the information gathered with a particular sensor. (ii) the rate of convergence and the stability of our mapping efforts in the face of typical uncertainties in navigation and attitude; (iii) as a methodology for actual sensor deployment and use on a real vehicle; and (iv) in tasks such as post-mission analysis for applications such as change detection and path planning for subsequent missions.

Thesis Supervisor: William Kenneth Stewart, Ph.D.

Title: Associate Scientist, Woods Hole Oceanographic Institution

Acknowledgements

For helping me through five and a half years of graduate school I would like to first and foremost thank my Official Windsurfing Sponsor.

I'd also like to thank:

The number of grad students who shared the joy, the misery, the windsurfing, the cheesecakes and the office space - Jamie Anderson, Todd Morrison, Franz Hover, Ram Gopalkrishnan, Gopal Rajagopal, Tad Snow, and Helen Huang.

The large number of U.S. Navy officers who came and went through DSL and with whom I interacted - Curt Murphy, Erik Burian, Gary Edwards, Larry Galvin, Rob Keefe, Ken Malmquist, Tom Woodford.

All the people at DSL which made it fun: Bob Ballard, Andy Bowen, Dezhang Chu, Tom Crook, Bob Elder, Larry Flick, Beven Grant, Mark Grosenbaugh, Steve Gegg, Skip Gleason, Linda Lucier, Anita Norton, Kathy Offinger, Will Sellers, Cindy Sullivan, Louis Whitcomb.

The Imaging group:

- Sysjan Marty Marra: for wearing his glasses backwards and the keyboard story.
- Jon Howland: for encouraging me to run and introducing us all to The Polar Bear.
- Steve Lerner: for teaching me structured programming practises.
- Dan Potter: for the critiques.
- Diane Di Massa: for the cheesecakes and advice.
- Xiaou Tang: for putting up with my concerted efforts to teach him C.
- The Emperor Min Jiang:
- John Kusters: for the fun classes we took together.
- Dan Leader: for his Kalman Filter and for waking me up early.

The ABE group: Rod Catanach, Al Duester, Steve Liberatore.

Ralfie for showing me how to sink a boat and raise the level of the ocean - all in one afternoon.

And the members of my committee (both official and unofficial): Al Bradley for brainwashing

me about the importance of energy, Rob Fricke for keeping my integrals straight, Eric Grimson for his astute comments, Nick Patrikalakis for being there, and my chairman Dana Yoerger who taught me how to debug code at sea.

My parents: Brijbala Singh and Pramod Singh.

Thanks to the Office of Naval Technology which paid for it all. This work was sponsored by ONT Contract# N00014-90-C-0098.

And lastly: The Great Communicator, The Computer Fairy, The Amazingly Thorough Editor, Santa Claus — Ken Stewart. Thanks Ken for giving me the pleasure of working at DSL.

Contents

Abstract	3
Acknowledgements	4
List of Figures	8
CHAPTER 1 Introduction	10
1.1 The Case for AUVs	11
1.2 Technologies for Underwater Imaging	15
1.2.1 Measures of Imaging Performance	15
1.2.2 Optical Imaging	16
1.2.2 Acoustic Imaging	17
1.2.2.1 Historical Survey of Imaging Sonars	17
1.2.2.2 Sidescan Sonar Technology Today	18
1.3 Sonar Sensing for Land-Based Robots	21
1.4 A Survey of Existing AUVs	23
1.5 Sensing for AUVs	23
1.6 Organization of this Thesis	25
CHAPTER 2 Stochastic Modelling for Sonars on AUVs	26
2.1 AUV Constraints on Imaging	26
2.1.1 Power Considerations	27
2.1.2 Vehicle Hydrodynamics	29
2.1.3 Navigation and Attitude	29
2.1.4 Computation and Telemetry Requirements	29
2.2 Typical Sonar Imaging applications from an ROV or AUV	30
2.3 The Sonar Equation and Stochastic Modelling	33
2.4 Stochastic Modelling: Some Elaborations and Extensions	35
2.5 The Effect of Navigation and Attitudinal Uncertainties	37
2.5.1 Uncertainties in Navigation	38
2.5.2 Uncertainties in Attitude	43
2.5.3 Computational Issues	45
CHAPTER 3 An Entropic Formulation for Stochastic Modelling	46
3.1 The Need for an Efficiency Measure	46
3.1.1 Desired Features for an Efficiency Measure	47
3.2 Entropy	49
3.3 Entropy for Sonars	52
3.4 Suitability of an Entropic Efficiency Measure	54
3.4.1 Incrementally Measuring Information Content	54
3.4.2 Multisensor Sonar Modelling	54
3.4.3 The Effect of Linear Operators	55
3.4.4 A Measure of Sensor Efficiency in a particular Environment	57

3.4.5 Consistently Assimilating Conflicting Information	60
3.5 Entropies of Emptiness and Occupancy	63
3.6 An Independent Validation: The 10% Rule	64
CHAPTER 4 Modelling Sensor Deployment	67
4.1 Entropy Change for Repeated Pings	67
4.2 Effect of Speed on Sensor Efficiency	68
4.3 The Effect of Navigation Uncertainty	72
4.4 Optimizing the Overall Efficiency of Mapping	73
CHAPTER 5 Application Data Sets	75
5.1 Shallow-water Mapping	75
5.1.1 Imagenex Sensor Characteristics	76
5.1.2 Navigation and Attitude	76
5.2 Computational Issues in 3D Mapping	81
5.3 Mapping Results	84
5.4 Entropy as a Measure of Stability and Convergence	89
5.5 Acoustic Mapping with ABE	95
CHAPTER 6	100
6.1 Post Mission Analysis	100
6.2 Analysis of Convergent Map	102
CHAPTER 7 Conclusions	105
7.1 Contributions of this Thesis	105
7.2 Directions for Future Work	106
APPENDIX A: Derivation of the Range Equation	109
APPENDIX B: Kalman Filtering for Underwater Vehicles	111
B.1 Kalman Filter Formulation	111
B.2 Navigation and Attitude Sensors	113
B.2.1 Navigation	113
B.2.2 Attitude	115
B.3 Heading Reference	117
APPENDIX C: Geometric Modelling for 3D Range Sensing	120
References	124

List of Figures

1.1 Mission Characteristics for the Autonomous Benthic Explorer	14
2.1 Mapping from an AUV: Sub-system Block Diagram	27
2.2 The JASON Remotely Operated Vehicle	32
2.3 Detectivity for 0.5" Circular Transducer	37
2.4 Probability Sensor Cone for Detectivity shown in fig 2.3	37
2.5 Original Sensor Cone	42
2.6 Sensor Cone for Navigational Uncertainty = 2m	42
2.7 Along Axis Cross-sections of Sensor Cones with Navigational Uncertainty	43
2.8 Across Axis Cross-sections of Sensor Cones with Attitudinal Uncertainty	44
3.1 Interacting Modalities for Mapping, Simulation and Motion Planning	47
3.2 A General Communications System	50
3.3 Entropy for a Bi-valued Distribution	55
3.4 Entropically Measuring Navigational Uncertainty	59
3.5 Entropically Measuring Attitudinal Uncertainty	59
3.6 Entropic Attitudinal Error for Occupancy	60
3.7 The Dempster-Shafer Rule of Combination modified for Stochastic Modelling	63
3.8 Entropy of Occupancy as a function of Altitude for a Given Swath Width	66
4.1 Range vs Speed for Different Hotel Loads	68
4.2 Entropy Change for Repeated Pings	69
4.3 Simulating Different Speeds for Optimizing Sensor Efficiency	70
4.4 Normalized Entropy Change vs Speed: Pencil-Beam Sonar	71
4.5 Normalized Entropy Change vs Speed: Broad-Beam Echo Sounder	71
4.6 Normalized Entropy Change vs Speed: Pencil-Beam Sonar w/ Nav Error	72
4.7 Normalized Entropy Change vs Speed: Broad-Beam Sonar w/ Nav Error	73
5.1 JASON Navigation for the Shallow Water Survey Operation	78
5.2 JASON Depth for Shallow Water Mapping	79
5.3 JASON Attitude (degrees) vs Time (samples at 2Hz)	80
5.4 Precomputed Overlapping Sensor cones	82
5.5 Variation in Phi for the Sonar Sensor Beam	82

5.6 Histogram of Returned Ranges	84
5.7 Aerial View of the WHOI Dock	86
5.8 200-kHz Side-Scan Sonar Data (Mechanically Scanned)	87
5.9 3-D 675-kHz Imagenex Data	88
5.10 Change of Entropy versus Time	89
5.11 Volumetric Slices for the Stochastic Volume	90
5.12 Depth for Pseudo 2-d Mapping	91
5.13 Change of Entropy for Pseudo 2D Mapping	92
5.14 Change of Entropy vs Time	92
5.15 Pseudo 2D Mapping - 400 pings	93
5.16 Pseudo 2D Mapping - 800 pings	94
5.17 Pseudo 2D Mapping - 1200 pings	94
5.18 Pseudo 2D Mapping - 1600 pings	95
5.19 Geometrically Derived Depth Map for ABE at Herring Pond	97
5.20 Change of Entropy for Depth Mapping at Herring Pond	98
5.21 Energy vs Time for Depth Mapping	98
5.22 Instantaneous Energy per bit	99
6.1 Post Mission Analysis	101
6.2 Entropy for Convergent Map	104
B.1 Variance in Kalman Filtered Navigation	114
B.2 Pitch / Roll for Vehicle at Rest	116
B.3 Gyro Bias vs Time	119
C.1 Geometrically Gridded Data	121

CHAPTER 1

Introduction

This thesis examines certain issues underlying underwater mapping from an Autonomous Underwater Vehicle (AUV). We begin by quoting a description of the work of the *HMS Challenger* [4], which sailed from Portsmouth in 1872:

With an inventory of 144 miles of sounding rope, 12.5 miles of sounding wire and other instrumentation the civil and naval crews... measured the total depth of the water... They found they had to make their depth measurements with hemp rope because the wire tended to kink and break.

Moving to 1986, more than a century later, we quote from a description by Dr. Robert Ballard [5] as he approached the *Titanic* in the submersible DSV Alvin:

A sight unlike any (we) had encountered in hundreds of deep ocean dives: directly in front of us was an apparently endless slab of black steel rising out of the bottom - the massive hull of the Titanic... But no sooner had I persuaded Ralph to let me look through his forward viewport, than he dropped Alvin's weights and we began hurtling toward the surface.... For less than two minutes we had glimpsed Titanic's towering hull. That was all we had to show for six hours of work.

And as one last sampling, we quote a geologist, Dr. Susan Humphries, from some of her work with the DSL120 and ARGO towed vehicles at a hydrothermal vent site in 1994:

Although most of the high temperature black smoker activity is confined to the summit of a small cone (~20 m in diameter) in the NW segment of the mound, several isolated black smokers have also been identified on the upper terrace. In addition, less vigorous black smoker fluids emanate from small rift parallel (NNE) fissures. A 20-30 m wide depression about 10 m deep, which appears

to be a recent collapse feature, is located immediately east of the black smoker cone, and a black smoker chimney is now being constructed on the lip of this depression.

These three glimpses into deep ocean exploration reflect some interesting trends in oceanography over the period they cover, while highlighting the strengths of the various technologies each depicts. Much has changed since the time of the *Challenger*. Perhaps not as obvious are the subtle differences in using a manned submersible as opposed to a towed vehicle. The submersible serves to immerse the participant in the underwater environment but, of necessity, the overriding concern in the use of submersible technology is for the safety of the occupants. On the other hand, the towed vehicle, with its large suite of sensors provides much more quantitative results, albeit at the cost of removing the scientist from the immediacy and excitement of being at the bottom of the ocean.

These trends reflect the increased emphasis on obtaining quantitative information in oceanography; it being no longer acceptable to simply make qualitative observations.

However, the following verse from Punch, sending the *Challenger* on its way, still holds:

*Her task's to sound Ocean, smooth humours or rough in,
To examine old Nep's deep-sea bed....
In a word, all her secrets from Nature to wheedle,
And the great freight of facts homeward bear.*

In fact, one could use this verse to describe the AUVs being conceived to work in the ocean today and well into the decades to follow.

1.1 The Case for AUVs

Before embarking on a course of study to qualify the methodologies of underwater mapping with

AUVs, one might question the fundamental need for such vehicles. Proven technologies associated with manned submersibles, towed vehicles, and ROVs already exist. These are, and have been, delivering the kind of quantitative results that scientists require for their work in the oceans. We try to answer the questions about the needs for AUVs by briefly examining the technologies associated with manned submersibles and Remotely Operated Vehicles (ROVs).

Manned submersibles, which were operational more than twenty-five years ago, are a viable technology even today. Despite limitations of power and severe constraints imposed on their design and operation because of concerns about operator safety, these are the workhorses of the deep ocean for tasks involving manipulation. They provide an intense experience of the deep ocean to the people in the submersible, even though they may be limited in the amount of information they can gather in comparison to other available technologies.

Tethered vehicles, both towed and ROVs, have an ample supply of power, very high-bandwidth real-time links to the surface, and a plethora of available optical and acoustic sensors. They can deluge scientists with vast amounts of data and are good survey tools, and provide the ability to react in real-time to opportunities that may present themselves. But these advantages come at the price of increasing complexity. The pilot also has to fly the vehicle from on-board ship with the presence of a tether making the task even more complicated.

Both these technologies are expensive and unsuited for protracted operations at a site. Each requires a ship to be in attendance at all times. The only extant option for long-term studies at a particular site is the deployment of moored arrays, but these are limited to making point measurements. There exists among the scientific community a need to perform long-term spatiotemporal studies cheaply and effectively.

For example, scientists are very interested in understanding the time scales of change that occur

in the various geological and biological processes at a hydrothermal vent area. As illustrated in Figure 1.1, which outlines the mission characteristics of the Autonomous Benthic Explorer (ABE), this is a niche (among others) — long-term spatiotemporal sampling of a remote site — that AUVs are trying to fill.

There are a host of other possible application areas. These include monitoring a remote hazardous area such as a toxic dump site; undertaking hazardous tasks such as those associated with mine counter measures; and working in inaccessible areas such as under the ice in the Arctic ocean.

At this juncture, one should warn against raising false hopes about the capabilities of AUVs. Achieving autonomous behavior in its basic forms is very hard to realize. The truth remains that in the near term AUVs will not be in a position to replace the high bandwidth acoustic, video and still camera work being carried out by the current generation of ROVs or, for that matter, the manipulative abilities that are the strength of manned submersibles.

It is interesting to compare the views of the deep ocean community, which realizes the importance of manned and unmanned vehicles, with the views of the space community which is grappling with some of the same issues. Despite shrinking budgets, it is a testament to the foresight of the oceanographic community that we have decided to pursue each of these technologies rather than being forced to choose one over the other as in the case of the space station. And this difficult choice has been made, even though we are operating at an infinitesimal fraction of the total budget of the space community.




Figure 1.1 Mission Characteristics for the Autonomous Benthic Explorer (ABE).

1.2 Technologies for Underwater Imaging

After making a case for AUVs, we focus our attention on understanding the capabilities and limitations of imaging from AUVs by examining the two dominant imaging modalities for working underwater - optical and acoustic. We realize immediately that there are numerous trade-offs involved in deciding to use a particular imaging sensor. A good review of the issues involved in imaging underwater and the underwater platforms from which imaging sensors are frequently deployed can be found in two related articles by Stewart [104] and Yoerger [125].

1.2.1 Measures of Imaging Performance

We digress here briefly, to talk about basic measures of imaging performance. Judging the quality of an image is a very subjective process. In some isolated cases, as in mapping with a range-based sensor, the quality of an image can be assessed directly by comparing the map so constructed, with independent ground-truth information. But the majority of imaging processes produce images where the image quality is subjective.

Besides the nature of the image, there are other considerations in evaluating a particular imaging modality. We must also consider the robustness, accuracy, and computational complexity associated with using the particular modality. While we would like superior performance with respect to each of these measures, the requirements for the individual measures are often in direct conflict with one another, and our specific realization of the imaging modality must make choices and trade-offs to optimize for the task at hand. Throughout the rest of this thesis, these issues, whether explicitly mentioned or implied, are among the basic constraints that lead to a particular algorithm or formulation.

1.2.2 Optical Imaging

Let us now consider optical sensors. The rapid attenuation of electromagnetic radiation underwater limits the use of such sensors to ranges of a few meters in the best of circumstances. A recent advance in laser technology has extended this range to tens of meters but the fact remains that the use of optical sensors is limited underwater because of range considerations. However, despite this severe limitation, when we look at sensor deployments in the field, video and CCD cameras still outnumber acoustic devices. The preponderance of optical sensors is partly a function of the fact that expertise in optical imaging technology can be transferred, even if only in a limited manner, from existing land-based applications to the underwater environs.

There are major differences in optical imaging underwater as opposed to optical imaging on land. With the power constraints and the physical size constraints associated with underwater platforms, it is not immediately obvious how one might go about achieving an optimum lighting geometry for a particular task on a particular platform using a particular sensor. Basic requirements, such as ensuring separation between the light source and the camera to minimize backscatter, are easy to formulate, but the task of underwater lighting, despite some progress in the field [61], remains a black art.

Work has also focussed on using lasers to derive profiles of the seabottom [68][65][123] as well as to provide scale information, range information, orientation, area of coverage, and details about feature size [114].

Research efforts in optical imaging underwater includes taking concepts already proved on land, such as stereo vision [52], shape from shading [56], and automated image segmentation and classification, and adapting them to the underwater environment [104][85][35]. Efforts are also

focussed on inherently underwater applications such as photomosaicking with underwater vehicles [77], correlating acoustic and video imagery [103], and image-compression technologies for acoustic telemetry [28].

There are fundamental problems with only using optical cameras on an AUV. A typical operating scenario would force the AUV to work close to the bottom to obtain good imagery. Algorithms for the automated interpretation of this imagery, in the absence of a human operator, are either computationally very expensive or not robust enough from the standpoint of vehicle operational integrity and force us to look for an alternative means of path planning and obstacle avoidance.

1.2.2 Acoustic Imaging

Given the nature of underwater optical imaging and its inability to cope with the basic tasks of path planning and obstacle avoidance for AUVs, we now turn our attention to acoustic imaging.

1.2.2.1 Historical Survey of Imaging Sonars

An interesting description of the evolution of acoustics is provided by Urick [116]. He traces the science of acoustics from first attempts in the early nineteenth century to measure the speed of sound in water to the development of echo sounders for detecting underwater objects. He goes on to describe the impetus of the science through the two World Wars and beyond.

For the purpose of underwater imaging, Russell-Cargill [95] points out that work in the second World War led to the realization that at low frequencies sound could penetrate the seafloor, which eventually resulted in the development of sub-bottom profilers [48]. High-frequency reflections from the seabed led to the discovery of the correlation between returned signal intensity and the

material that composed the seafloor, which led to the development of the first side-scan sonar at the National Institute of Oceanography (NIO) in England [111].

In the United Kingdom, the NIO, now known as the Institute of Oceanographic Studies (IOS), continued work on side-scan sonar, resulting in the development and use of the Geological Long-Range Inclined Asdic (GLORIA) [94] system and its successor the GLORIA II [101].

In the US, the development of side-scan sonar was carried out by a number of scientific institutions as well as by industry. Notable examples include the Klein family of sonars developed by Klein Associates [64], the SeaMarc systems [14] developed jointly by University of Hawaii and International Submarine Technology Limited and the SeaBeam systems developed by General Instruments, later SeaBeam Incorporated.

1.2.2.2 Side scan sonar technology today

Recent developments in the field include the use of phase-difference techniques on calibrated systems, such as the SeaMarc II and the DSL120. These collect coregistered bathymetry and calibrated intensity in order to do an automated analysis for the segmentation and classification of bottom types. Research in the field has also focussed on understanding the geometrical and environmental effects on the acoustic sensing modality.

Acoustic imaging devices are all governed by the Sonar Equation [44][116]. The nature of acoustic imaging underwater is such that the returned intensity for a particular image is a combination of quantities linked to the physics of the acoustic sensing modality and the geometry of the environment. The various generations of sensing systems have incrementally been building up a repertoire of processing mechanisms to deal with various sources of error. We shall revisit in some

depth in Chapter 2, some of these issues for the sensors dealt with in this thesis, but the reader may also wish to refer to specific examples of processing side-scan sonar data from GLORIA [79], Sea Marc [90][91], Sea Beam [36][37] and the DSL120 [107][109] for concrete examples of methods of processing sonar data for the effects of beam patterns, transmission losses, and geometric and radiometric effects such as slant range correction, the removal of water column data, and segmentation according to material type.

Considerable work is also being carried out to understand the basic interactions of high-frequency acoustic waves with the seafloor [9][29][43][51][60], the sea surface [45][117], and in shallow water[106][58]. Another development, fuelled by the needs of the offshore oil industry, has been the advent of high-frequency profiling sonars. Current manufacturers of such systems include Simrad, Tritech Systems, Reson and Imagenex.

Table 1 summarizes several sensors that exist today along with important considerations for their use and deployment in the field. Here, LBL stands for low-frequency long baseline systems [81]; EXACT [101] is a commercially available high precision system. We note here that the large number of successful acoustic mapping operations have, with notable exceptions [105], all been carried out with towed or ship-mounted sensors. The difficulty in acoustic mapping with ROVs and AUVs lies not in the actual sensors but rather in integrating these sensors on the platform in question. When the accuracy of our navigation system is lesser than the accuracy of the resolution of the sensor by orders of magnitude then this shows up as inconsistencies in the mapping efforts when we try to grid the sensor data to the limits of sensor resolution. There are also fundamental concerns about sampling the grid at a fine enough resolution.

Table 1: Sensors and their Deployment Characteristics

Sensor	Resolution, Order of Magnitude (m)	Range	Platform	Nav System (accuracy in m)	Nav Accuracy / Resolution Ratio
Laser	$10^{-3} / 10^{-2}$	tens of meters	ROV, AUV, Towed	Exact, LBL ($10^{-1}, 10^0$)	$10, 10^2$
CCD Cam- era	$10^{-3} / 10^{-2}$	tens of meters	ROV, AUV, Towed	Exact, LBL ($10^{-1}, 10^0$)	$10, 10^2$
675kHz Mesotech	10^{-2}	several tens of meters	ROV, AUV, Towed	Exact, LBL ($10^{-1}, 10^0$)	$10, 10^2$
DSL 200 kHz	$10^{-2} / 10^{-1}$	hundred meters	ROV, AUV, Towed	Exact, LBL ($10^{-2}, 10^0$)	1, 10
DSL 120 kHz	10^{-1}	up to 1 km	Towed	LBL (10^0)	10
HMR1	$10^1 / 10^2$	several km	Surface Ship (Towed)	GPS (10^0)	10^{-2}
SeaBeam	10^2	several km	Surface Ship	GPS (10^0)	10^{-2}

Table 1 above offers a qualitative explanation for the problems with gridding high-frequency sonar data with typical navigation systems. If an acoustic sensor is hull mounted on a ship one can take advantage of all the instrumentation available on board along with the luxury of GPS navigation. When we consider the accuracy of navigation (a few meters), relative to the resolution of the sensor along the bottom (a few hundred meters), the errors turn out to be relatively small. For towed vehicles, the navigation accuracy to the resolution of the sensor ratio is larger but still acceptable. For ROV (and even more so for AUV) operations, the navigation accuracy to the resolution of the sensor ratio grows such that, coupled with the noisier operating environment, limited space, limited weight, limited power, and matters of expense, we wind up being severely handi-

capped in our mapping efforts.

1.3 Sonar sensing for land-based robots

Let us step back at this juncture and examine the experiences of a parallel stream of thought, the land-based mobile robot community. The problem of perception for autonomous land-based mobile robots is still an open issue. Some encouraging results have been reported [2][3][71] using optical sensors, both video and laser, but sonar as a sensing modality is one whose potential has only been partially realized.

It should be pointed out that the land-based community looked towards sonars for entirely different reasons than those associated with the oceanographic community. Optical sensors are not limited in their range or resolution on land as compared to acoustic sensors. The reason behind the interest in acoustic sensing for land-based applications had more to do with the price of these sensors. Acoustic sensors held out a promise of providing a very cheap ranging system. In using an acoustic sensor, the land-based mobile robot community has primarily used the Polaroid transducer [88]. The simple time-of-flight measurements associated with these sensors seem to provide a cheap and easy method for sensing the external environment. The reality, unfortunately, is that the Polaroid sensor is deficient in a number of ways. Some of the major sources of error [73] are differences in the duration of the transmitted pulse, a poor 16-step time-varying gain, and errors in the threshold-detection circuit (which in turn stem from the method of capacitive charge-up that has been physically implemented).

Notwithstanding these shortcomings, much effort has been put into using the Polaroid sensor for mobile-robot perception. A number of methods were used to overcome what was thought to be the “noisy” nature of these transducers. Moravec and Elfes [47][83] introduced the concept of proba-

bility grids to describe the uncertainty associated with sensor readings and to combine multiple, possibly inconsistent, readings. Drumheller [42] used the Polaroid sensor for localization. He describes an algorithm that exploits geometric constraints between the sensed data and an a priori map to obtain the position and orientation of the robot within the map.

While these initial approaches spelled out the potential for using sonar, researchers were unable to extend or elaborate on these theories. The consensus among the mobile robot community has been that sonar is not a very useful sensor [22][33][42].

The bad reputation acquired by sonar has been due to a combination of factors. For a start there is the virtual synonymy of the Polaroid transducer with the sonar sensing modality for land-based robots. This synonymy has led researchers to ascribe all the bad points of the Polaroid sensor to the sonar sensing modality in general. Also, there has been, until quite recently, a lack of appreciation for the physics behind acoustic sensing. Multipath and radiometric (due to different material types) effects, when taken in combination with the nature of the Polaroid sensor, precluded any large measure of success.

Recently, in pursuing a global mapping algorithm Moravec [82] has been working on basis functions that are automatically learned by sensing the environment. It appears that, given the physics of acoustic sensing, with its interplay of quantities of very large dynamic range and variability, the successes of such an approach will be limited, if any, to a very constrained environment.

There has also been a renewed interest in using Polaroid sensors for sonar-based localization. Kuc and Siegel [67] pointed out differences between reflecting surfaces and diffracting objects and derived target models for corner, wall, and edge targets. Leonard [73] defined so-called “regions of constant depth” and used an extended Kalman filter for model-based localization using the Polaroid sensor. Part of his success lies in the fact that he realized and tried to compensate for the

various shortcomings associated with the Polaroid transducer. He also paid close attention to the physics of acoustic sensing in arriving at his regions of constant depth.

1.4 A Survey of Existing AUVs

Several AUVs have been developed and built over the years. Most of these were designed solely to follow a pre-programmed path in the water column with the emphasis of the work focussed on developing the control system and other vehicle propulsion considerations. The earliest vehicles conceived include the EAVE vehicles [15][16][17], designed by Blidberg et al.

A number of the current and past efforts are driven by military needs. These efforts have included the XP-21 vehicle [129] developed by Applied Remote Technology, the UUV vehicles developed at Draper Laboratories [129], and the LDUUV vehicle [129] developed at the Naval Underwater Weapons Center.

AUVs being proposed for scientific use in the United States include the OdysseyII [12] which is slated to be the vehicle to be used with the Autonomous Ocean Sampling Networks, the Autonomous Benthic Explorer [127] and the Remote Environmental Measuring Units [118].

Abroad, the standouts include the Norwegian Defence Research Establishment's NDRE-AUV [62], the Marius vehicle [46] being designed at the Instituto Superior Tecnico in Portugal, and the Ptero A150 [115], and the Aqua Explorer [66] developed in Japan by the University of Tokyo and Kokusai Denshin Denwa Company.

1.5 Sensing for AUVs

In the preceding sections we have shown that we require some form of acoustic imaging for

AUVs to provide the basic support functions associated with obstacle avoidance and path planning. We also made a case for using acoustic imaging for underwater mapping. Thus, when we consider acoustic imaging for AUVs, it seems natural to merge two distinct streams of thought already alluded to in previous discussions. The first, arises from the experiences of the underwater imaging community and the second from the research of the land-based robot community, which has been working on perception issues for obstacle avoidance and path planning for a long time.

Before attempting to merge these two different paradigms, however, it is necessary to look at the constraints particular to AUVs. In some ways constraints on AUVs are much more restrictive than those for land-based robots. Walsh [120] provides a good summary of some of the issues involved in the design and use of AUVs.

The ocean is a hostile environment to work in. Typically, too, the regions of operation are the harshest of the hostile. For example, the Autonomous Benthic Explorer has been designed to work in hydrothermal-vent areas; Odyssey's missions have been under the ice, in the Arctic Ocean, and the LDUUV has been designed to work on mine countermeasure operations in shallow water.

A major constraint is the limited amount of energy available. If we consider typical numbers for ABE [126], it becomes obvious that the dominant energy requirement on board a small AUV is the power requirement for sensors. The tens of Watts required for a typical acoustic or optical sensor easily matches the sum total of power requirements for propulsion and computation for the rest of the vehicle.

Another constraint arises from the very low bandwidth available between the vehicle in the water and its operator on the surface. Considerable progress has been made in this direction recently [28], but we are still not at the stage where we can do real-time telemetry of all the incoming sensor data without some scheme to sift out the "interesting" data.

Yet another limitation of imaging from an AUV is the fact that vehicle designers choose very efficient torpedo shaped profiles for their vehicles based on power considerations aimed at minimizing drag. Unfortunately these profiles exhibit poor attitudinal stability, which directly affects their imaging performances.

1.6 Organization of this thesis

With this background we now turn our attention to the remainder of this thesis which presents a formulation to analyze the issues discussed in this introductory chapter. Chapter 2 looks at fine scale sonar mapping underwater within a stochastic framework. We present a strategy for developing a stochastic model for a sensor and show how this stochastic model may be extended to deal with navigation and attitudinal uncertainty. Chapter 3 proposes and justifies the formulation of an entropic measure of efficiency for our mapping efforts to be used with the stochastic model outlined in Chapter 2. Chapter 4 looks at a series of simulations to examine the use of our formulation for mission-level activities such as optimizing the speed of the vehicle and the effect of the beam pattern on our mapping strategies. Chapter 5 examines the application of an entropic formulation to data acquired using the Jason ROV. We show the stability and rate of convergence in challenging shallow water environment and the applicability of our methodology to change detection in its environs. We also look at data acquired using the ABE AUV and the applicability of our methodology to short-term optimization strategies for imaging. Chapter 6 uses the stable convergent map and analyzes it to determine regions of interest for path planning activities. Finally, in Chapter 7 we offer some concluding thoughts regarding the applicability of our techniques and suggestions for future work.

CHAPTER 2

Stochastic Modelling for Sonars on AUVs

Let us now concern ourselves with formulating a framework for underwater mapping by an AUV. We repeat here a reference to the work of Stewart [105]. His stochastic modelling approach has been one of the few successes in the area of very fine-scale underwater mapping at high frequencies. One other method [80] shows promise, but so far has been limited to the artificial environs of laboratory settings and is further limited by the navigational accuracy (~ 1 cm) it assumes.

There have been two somewhat opposing schools of thought in formulating strategies for robot sensing. The global mapping school, currently in disfavor among the majority of the land-based mobile robot research community, is of the opinion that we should try to maintain an accurate, usually grid-based, map of our surroundings linked to the outside world. The local school, on the other hand, maintains a much simpler model of its environment centered around the robot.

While one could argue the pros and cons of each of these schools, the reality is that successful autonomous robots will probably utilize some combination of both these methods. For the purposes of this thesis, it may be pointed out that the formulation of imaging presented in this and succeeding chapters is independent of and equally applicable to both these methodologies.

2.1 AUV Constraints on Imaging

We look first at the constraints that arise in working with AUVs. The importance of integrating hardware and software for a systems-based approach to computing is widely recognized, and we hope in this section to underscore the importance of integrating imaging strategies with the design

of AUVs. Figure 2.1 shows some of the relationships between various sub-systems on an AUV and the imaging sensors deployed on it.

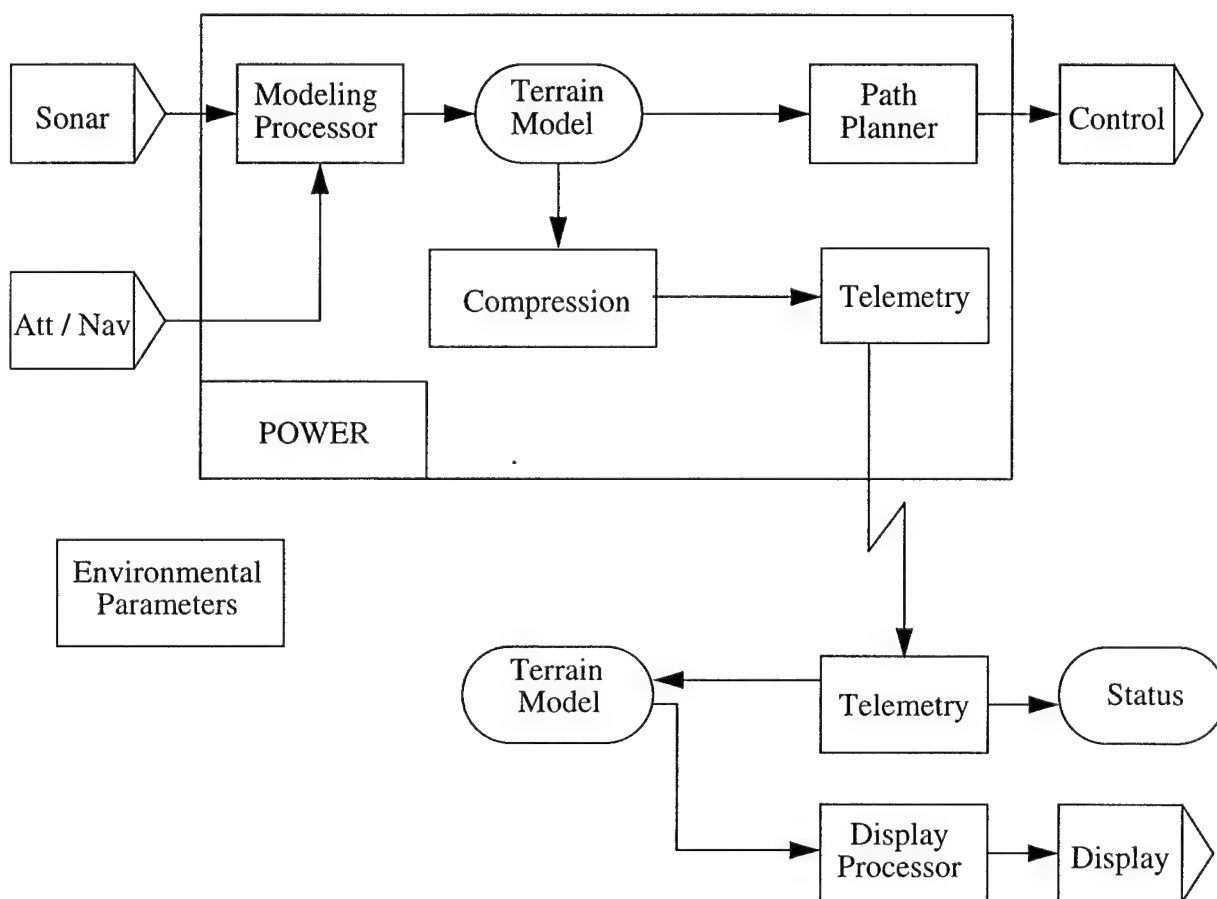


Figure 2.1 Mapping from an AUV: Sub-system Block Diagram

2.1.1 Power Considerations

The most constraining issue facing AUVs today is power. The oceanographic community has always been interested in efficient ways of storing energy for long-term equipment deployment. An anecdote related by a colleague [122] in the oceanographic community comes to mind. He was

charged with watching the journals for technical breakthroughs in energy storage. In years of watching the journals, the only publication of significance that he could report described work in reducing the thickness of the jacket enclosing standard cells. This stratagem allowed the authors to vastly increase the amount of electrolyte and hence the energy storage capacity of the cell.

The truth remains that battery technology will not be revolutionized anytime soon. Most AUVs use ordinary commercially available lead-acid batteries. Other esoteric technologies are in various stages of development or available at considerable expense. These technologies include silver-zinc cells, lithium cells, and aluminum-oxygen semi-fuel cells. Trade-offs in using each include issues of cost, handling, and efficiency. Suffice to say that the kinds of missions on the drawing board for AUVs today will easily be a challenge to the technologies of energy storage for some time to come.

It behooves us, then, to look to our designs and imaging technologies to tighten the belt on vehicle energy consumption. Typical numbers for propulsion are several tens of Watts, and the value is roughly of the same order of magnitude, if not slightly higher, for acoustic imaging sensors. The vehicles are already incredibly energy efficient. Oceanographic instrumentation has a rich tradition in efficient circuit design for long-term deployments, and this has yielded some very energy conscious AUV designs.

The other option is to look at acoustic imaging transducers. Most sonar sensors achieve 10-50% efficiency in converting the incoming electrical energy to outgoing acoustic energy, and there are no new technologies on the horizon to enhance these efficiency ratios. Besides, even at 100% efficiency, we must put forth significant amounts of energy (relative to the power budget of the vehicle) to physically ensonify the area being mapped.

The imperative, then, is to try to fully understand the value of the imaging data in a given environment so that we can justify our power consumption from a mission-level standpoint.

2.1.2 Vehicle Hydrodynamics

Most AUVs being developed today are torpedo shaped, as this shape yields the lowest drag. Low drag in turn carries with it the benefits of energy reduction for propulsion of the vehicle. However, while reducing drag, hydrodynamics also greatly affect the imaging performance for a particular platform. The popular torpedo-shaped vehicles exhibit very poor inherent roll stability, which can have a serious effect on the performance of imaging sensors mounted on the vehicle. Options for fixing roll instabilities include fins and active control algorithms, but the fact remains that imaging and vehicle performance may sometimes be at odds with each other and experts in both fields need to interact closely to achieve overall mission goals.

2.1.3 Navigation and Attitude

A typical AUV mission, a long-term spatiotemporal study of a hydrothermal-vent area, for example, requires precise control algorithms for track following and sensing, which are in turn dependent on the availability of good on-board sensors for attitude and navigation. While a temporary loss or deterioration in the value of one of more of the navigational or attitudinal sensors may not compromise vehicle integrity, it does seriously impact the value of any imaging being carried out. The requirements for good imaging are much more stringent than the requirements for control and possibly could require greater bandwidth.

2.1.4 Computation and Telemetry Requirements

In both these cases, the imaging requirements dominate the overall vehicle architecture. The most

computationally expensive task on board a vehicle is usually the control system, but this computational load will be dwarfed by the requirements associated with imaging, as vehicle attitude and navigation needs to be merged with incoming imaging sensor data to form some representation of the environment for the purposes of mission planning and obstacle avoidance.

From a telemetry standpoint, given the relatively low bandwidth available and the nature of the mission for a typical AUV, the vehicle itself would only be passing status information. However, for imaging purposes, we need the ability to transmit some representation of the imaging process, which requires orders of magnitude greater bandwidth than status information.

2.2 Typical Sonar Imaging applications from an ROV or AUV

The discussion in the previous section was not meant to portray the underwater imaging modalities as a power-hungry, hydrodynamic-shape-dictating, bandwidth-grabbing, navigation-attitude-requirements-dominating, computation monster.

Rather, the purpose in highlighting these constraints is to examine the effect of imaging performance as these quantities deteriorate from their optimum values to values determined by a particular AUV design. This will, in turn, help AUV designers appreciate the effect of design choices on imaging and mission-level goals.

Underwater mapping from an AUV, relative to mapping from ROVs, is severely constrained because of factors discussed in the previous section. Although ROVs in turn have their own constraints with respect to mapping underwater (as outlined in Chapter 1), the immediate goal for underwater mapping from AUVs is to try to autonomously achieve the level of mapping performance attainable by ROVs. We thus turn our attention to mapping with ROVs to better understand the limits of current underwater mapping efforts. Figure 2.2 shows a typical ROV equipped for an

underwater mapping application. The imaging sensor here is a side-scan sonar, one transducer of which is seen mounted on the port side of the vehicle. Depending upon the application, it may be mounted in a variety of configurations — looking forward in an obstacle avoidance mode, pointed downwards to collect bathymetric information, or tilted to the side, as in this case, to achieve an optimum imaging geometry while minimizing cross-talk between channels. The navigation and attitude sensors are in their own housings distributed across the vehicle. The measurements from individual sensors may be made in a local frame of reference and must be transformed to some convenient common frame of reference from where they may be incorporated into our mapping algorithm.

The sonar operates in an imaging mode by transmitting a pulse of short duration, waiting for a small period of time to allow the pulse to be propagated away from the immediate vicinity of the transducer, then switching on the receivers, which digitize the intensity of the returned signal to form an image for display and storage.

The intensity of the returned signal is a function of a number of quantities, as described by the Sonar Equation, of large dynamic range. The process of transforming the intensity values into a consistent grid is described in depth in the next section.




Figure 2.2 The JASON Remotely Operated Vehicle.

2.3 The Sonar Equation and Stochastic Modelling

Imaging and sidescan sonars, like all acoustic devices, are governed by the Sonar Equation[116]. Various forms of the sonar equation arise from considerations for particular sensors. One form is given in (2.1) below.

$$RL = SL - 2TL - 2\alpha r + TS - BP_t + BP_r + FG + NL \quad (2.1)$$

where RL = Received Level

SL = Source Level

TL = One way transmission loss

α = Absorption constant

r = Path length

TS = Target Strength

BP_t = Beam Pattern, Transmitting

BP_r = Beam Pattern, Receiving

FG = System Gains and

NL = Noise Level

An algorithm that works well for high-frequency acoustical imaging focuses on backprojecting probability values into a volumetric model. For calibrated systems each of the terms in (2.1) may be measured or otherwise accounted for. This simplifies the task of going from the sonar equation (2.1) to a probabilistic sensor cone. Unfortunately, for most commercial systems the individual values of these terms are not available (either because the terms have not been individually measured and / or because they are regarded as proprietary information). However, we can arrive at very

good estimates of the quantities in (2.1) for an arbitrary system in the following manner.

The beam patterns can be either measured in a test tank with a known source or they can be calculated analytically given the physical dimensions and the geometry of the transducer. For a circular transducer the analytical solution for the beam pattern is a Bessel function which approximates a conical beam in space. We further make the assumption that the transmitting and receiving beam patterns are identical.

The source level can be fixed by measuring the electrical energy being fed to the transducer and by assuming a reasonable value for the efficiency of the transduction process (we assumed a 10% and 40% efficiency for the broad-beam echo sounder and the pencil-beam sonar discussed in the later parts of this thesis). With a knowledge of the beam pattern we can now fix the source level at the transducer.

The transmission losses and the absorption constant can be easily fixed given the operating frequency of the sonar.

The terms corresponding to the target strength, noise level, and fixed gains are harder to fix. The target strength is a function of the grazing angle and the material type. Echo sounders are typically mounted in a geometry which ensures near normal incidence on the seafloor. The effect of differences in material type can lead to a change in target strength of the order of 20 to 30 dB. For mechanically-scanned pencil-beam sonars the grazing angle effects can add a 10 to 20 dB to the range of target strengths. The fixed gains in the system are an attempt to compensate for differences in target strength. The only conjecture we can make with any degree of certainty about the noise level is that most commercially available high-frequency sonar systems are limited by their self-generated electrical noise levels and that these levels are difficult to measure.

While the terms corresponding to target strength, noise level, and fixed gains are hard to quantify individually, we can estimate the sum of these terms taken as a whole by considering the maximum range of the system. The maximum range may be either obtained from system specifications or from experiments in the field. By substituting the value for each of the terms in (2.1) at the maximum range we can calculate the signal-to-noise ratio that is the minimum required for target detection. Note that even though we cannot fix the signal-to-noise ratio absolutely, we can use this value (which we denote hence by the symbol A) to calculate the difference between the signal-to-noise ratio at an arbitrary point as opposed to that at maximum range. This difference can be used to fix the probabilistic sensor cone that we wish to backproject into our stochastic model. The sensor cone that is eventually backprojected is thus generated by using a simplified model of the basic sonar equation (2.2) to calculate the term D , termed detectivity, which is used in combination with A as outlined below.

$$D = SL - 2TL - 2\alpha r - BP_t + BP_r \quad (2.2)$$

Step 1: Consider a set of volume elements, in a computationally efficient manner such that the set of elements considered is a superset of all the elements ensonified.

Step 2: For each of the elements identified above, calculate the range from the sensor and combine it with the location of the volume element to calculate the angle of incidence from the centre of the sensor beam.

Step 3: Calculate the detectivity, from (2.2), for each of the volume elements.

Step 4a: If the detectivity exceeds A , and the range of the volume element is less than the range returned by the sonar, assign a probability of emptiness via a mapping function f_l .

Step 4b: If the detectivity exceeds value A and the range of the volume element is within some value, say ϵ , of the range returned by the sonar, assign a probability of occupancy via a mapping function f_2 .

ϵ , f_1 and f_2 are determined based on sensor characteristics as explained in the next section.

2.4 Stochastic Modelling: Some Elaborations and Extensions

The utility of the algorithm presented above has already been shown in its applications to various datasets [105]. However, issues of stability and convergence of the algorithm and the choice of the key parameters that ensure consistency have not been resolved. These issues take on added importance in autonomous operations where there is no operator in the loop to filter out obvious or not-so-obvious noisy data points.

We now examine how we might assign a value to ϵ and determine the functions f_1 and f_2 .

ϵ may be fixed by examining the accuracy in resolution specified for the sensor.

Describing the functions f_1 and f_2 is more difficult. The choice of these functions is crucial to our algorithm. With the application of these functions, we map into a new domain where the operations we perform need not be consistent with the physics of the acoustic sensing modality.

We could take f_1 and f_2 to be some linear combination of the detectivity. They could just as easily be quadratic. In section 3.6 we provide an independent validation for our choice of f_1 and f_2 .

Before choosing, and attempting to justify our choices for f_1 and f_2 , let us examine other issues relevant to the algorithm. First, note that this method of assigning probabilities does not implicitly model target strength (TS) or the geometry of the object being imaged. We cannot make any a priori assumptions about either factor. The reality is that the object being imaged might have some arbi-

trary geometry and arbitrary target strength. We choose to ignore the effects of these two factors in our version of the sonar equation. For a particular ping, these two assumptions might seem as rather large leaps of faith, but in actuality our model is built up using a very large number of pings and these multiple pings help the volumetric model converge to the true geometry as will be shown later.

There are other factors to consider. An important issue is whether this approach is guaranteed to converge, especially in a typical AUV application where there is no human being in the loop. Further, even if the algorithm does converge, we would like to be able to characterize the rate of convergence to achieve higher level mission goals including the allocation of resources and energy to various sensors and other systems on the vehicle.

We choose the function f_1 to linearly scale the maximum value of D to 0.0 and so that A , the minimum acceptable value of D is mapped to 0.5. For f_2 we use a similar scheme except that we add a constant G to the value of the detectivity before remapping to the range of probabilities between 0.5 and 1.0. The choice for the value of G , which is a gain term that stands in for target strength, is limited to be less than A . However it must be large enough to allow for the detection of varying target strengths corresponding to different material types and grazing angles. We can arrive at a reasonable value for G by examining the backscatter curves for different material types, over the grazing angles and frequencies of interest.

We justify the choice of these functions by showing their utility and accuracy in the work of the succeeding chapters. One independent, though indirect, proof of their formulation is also given in section 3.6.

Figure 2.3 shows the values of D calculated for an ideal circular transducer of diameter 0.5" operating in the 200-kHz region after receiving a range return from 75m. Figure 2.4 shows the cor-

responding sensor cone for Fig 2.3. Notice that our stochastic sensor cone models the sidelobes. This might be a problem if we were to consider a single ping in isolation but our framework implicitly assumes that a particular region will be ensonified multiple times.

2.5 The effect of Navigation and Attitudinal Uncertainties

For typical ROV or AUV operations, accurate navigation and attitude are crucial for micro-mapping. We now consider each of these issues in turn.

2.5.1 Uncertainties in Navigation

Various options exist for AUV navigation. For long-term AUV deployment at full-ocean depths, the systems of choice are acoustically based. As with all acoustic systems, they trade-off range and resolution. Commercial long-baseline (LBL) systems typically provide accuracies of a few meters within an acoustic net spanning a few kilometers. Other systems, such as the EXACT, offer centimeter-level accuracy over regions up to 100m. While these systems are accurate enough to portray the navigation tracks of underwater vehicles and capture their dynamics, what often happens in practice is that the geometry of the net will lead to poor performance in certain regions and significant drop-outs in navigation may occur. While these drop-outs may not have much effect on the vehicle performance, which can sometimes get by in these regions by dead reckoning, they have a significant negative impact on mapping performance. Doppler logs of velocity are available on the commercial market with very good performance and at relatively low power. However, these are bulky, heavy, and somewhat expensive.




Figure 2.3 Detectivity for 0.5" Circular Transducer




Figure 2.4 Probability Sensor Cone for D shown in fig 2.3

How do these navigation uncertainties manifest themselves? To go back to an empirical argument, intuitively, what we are doing in mapping Detectivity to a probability, is estimating a target presence in proportion to the amount of energy going through the volume element it encompasses given the hypothetical noise characteristics at that voxel.

For a particular sensor in the presence of navigational uncertainties, Stewart points out that error and ambiguity are accounted for by blurring a spatial projection of the sensor data before back-projecting into the volume.

The process of spatially filtering the incoming data by using a blurring function is appealing to our intuition, for if we are uncertain about the whereabouts of the sensor we should correspondingly be uncertain about its sense of the external world. Also, stochastic backprojection is a process that assigns each voxel a probability of occupancy or emptiness corresponding to the amount of acoustic energy that passed through the voxel. If an uncertainty exists in navigation, the acoustic energy is redistributed according to our uncertainty, but in so being blurred the energy corresponding to a single ping cannot change. One possibility for such a blurring function is a unity gain, three-dimensional, boxcar filter in the time domain of spatial extent in each of the three axes corresponding to the uncertainty in that particular axis.

Figure 2.5 shows the result of using such a filter on the sensor cone shown in Fig 2.3 for a navigation error corresponding to 2m. Figure 2.6 shows the graphs along axis for the probability corresponding to sensor cones where we successively varied the error in navigation from 1m to 9m in 2m steps. The graphs show the peak of the probability of emptiness shifting away from the sensor. At first glance this seems counterintuitive, but on further reflection we can explain the trend in the following manner.

Close to the sensor, the beam is confined and the effect of the uncertainty in the navigation is

more pronounced. Further out, as the beam spreads out to values of the order of the navigational uncertainty, the effect of the navigation uncertainty is relatively diminished as the probability is dominated by the decrease in the Detectivity. Eventually, the the distribution falls back to its undisturbed trend.




Figure 2.5 Original Sensor Cone




Figure 2.6 Sensor Cone for Navigation Uncertainty = 2m




Figure 2.7 Along Axis Cross-Sections of Sensor Cones with Navigational Uncertainty

2.5.2 Uncertainties in Attitude

Cheap attitude sensors have become available for imaging platforms only in the last two decades. An interesting algorithm was put forward by Cobra [30], which compensates for roll, pitch, and heading errors by looking at the image statistics. Unfortunately, this algorithm is computationally expensive and applicable only to regions where there are distinct, man-made or naturally occurring features in the imagery.




Figure 2.8 Across-Axis Cross-Sections of Sensor Cones with Attitudinal Uncertainty

We can compensate geometrically for the attitudinal errors that we can measure but we may not be able to sample the attitude at a rate approaching the rate of our sensor. Further, there may be an error in the attitude measurement at the sensor level itself. In either case, if we are unable to fully capture the dynamics of the vehicle, we may introduce finite errors from simply upsampling our low-bandwidth attitude information. This in turn might have serious effects on the performance of our mapping algorithms.

We can proceed to treat the attitudinal errors in a manner analogous to our treatment of navigation errors, that is, by blurring the sensor cone to account for the uncertainty. For attitudinal uncertainties, however, we blur values in the appropriate angular sense. Figure 2.7 shows cross-sections

along the axis of the beam corresponding to errors in the attitude, which go from the original distribution to a 2.5° error in steps of 0.5° .

2.5.3 Computational Issues

Analytically and intuitively we may be quite satisfied with our results, but we have not really addressed the issue of real-time data processing. While spatial filtering to model the navigational uncertainties is computationally expensive, the real source of concern here is the radial filtering corresponding to uncertainties in attitude. Our underlying mapping representation is grid-based, and to accomplish radial filtering we need to sort the grid by radius, complete the filtering in this sorted domain, then unsort the results back into grid coordinates.

Using a spherical or cylindrical coordinate frame will not solve the problem, because in that case the awkwardness arises for our navigation uncertainty. We raise this point here to bring it into focus but put off a discussion of the solution until a later section.

CHAPTER 3

An Entropic Formulation for Stochastic Modelling

So far, we have extended the stochastic modelling framework to allow for navigational and attitudinal uncertainties. Even so, the framework has shortcomings that prevent its direct application to a real-world AUV. A human operator is required, for example, to filter out noisy data points. It goes without saying that “noisy data” could seriously harm one’s mapping efforts. The term “noisy data” itself is very subjective and hard to characterize. Human operators have the luxury of using their judgement in a particular environment as well as the option to reconsider their processing if the results do not match up to their expectations. We now concern ourselves with such a scenario for autonomous operations, namely, of quantifying the deterioration in mapping performance with the degradation in performance of the sensor modality.

3.1 The Need for an Efficiency Measure

We have already formulated an algorithm for treating uncertainties in navigation and attitude for a particular ping. But we would like to quantify the differences in some consistent manner for the different error values. Moreover, we would like to be able to compare the degradation caused by attitude errors with that caused by navigation errors.

While defining a measure of the efficiency for a sensor, it is instructive to catalogue and address some of the other issues related to the process of mapping from an AUV.

3.1.1 Desired features for an Efficiency Measure

1. It should be a consistent measure of the phenomena that constitute the sensing process.

While this seems an obvious point, it is nevertheless a point that must be made. Our measure of efficiency should apply to the different processes that make up the sonar sensing modality. The original motivation for deriving such a measure was to characterize the efficiency of different schemes for assimilating the incoming data. Additionally, we would like the measure to characterize the transduction process between the world and the sensor, the noise emanating from various sources within the sensing modality, and the efficiency of processing.

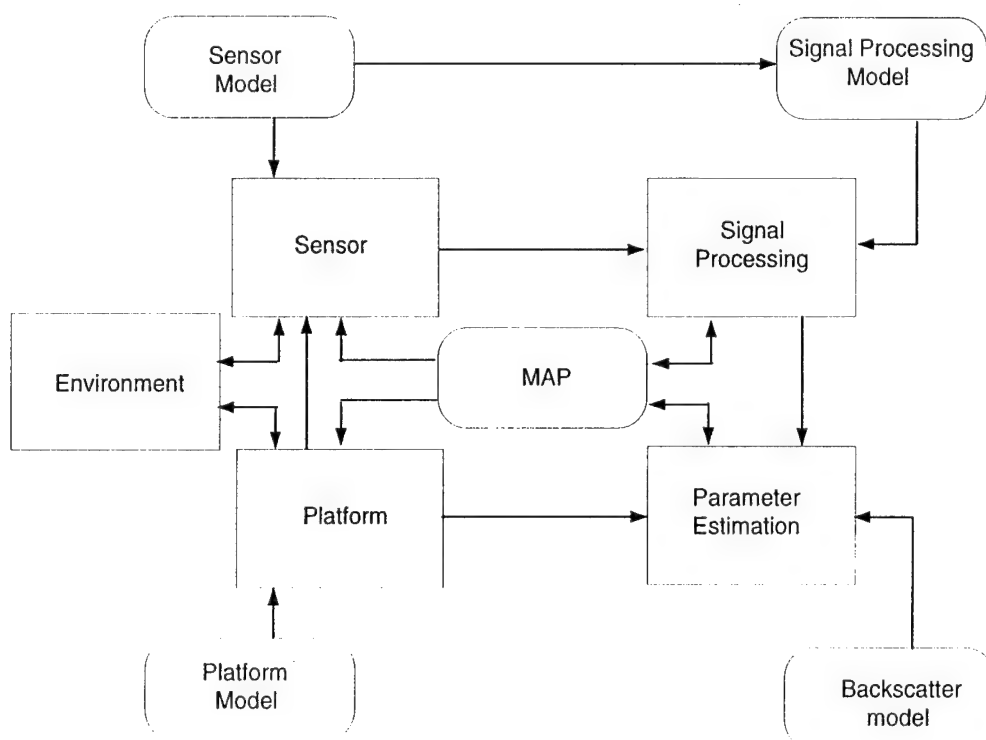


Figure 3.1 Interacting Modalities for Mapping, Simulation and Motion Planning

We point out here that such a measure need not be used exclusively for real-time mapping applications but also might find utility in simulation and mission-planning scenarios. As outlined in Fig

3.1, we should be able to characterize various aspects of these processes by the amount of information flowing from them and by how the information from a particular block affects our knowledge of the environment being modelled. Note that such a global efficiency measure need not exist. However, among different measures of efficiency, closeness to this characteristic is a desirable property.

2. The ability to consistently assimilate multiple values for a particular voxel.

In our sensing process, certain voxels might be interrogated several times while other voxels might not be interrogated at all. In building a stochastic model, a simple averaging process may be used to combine multiple probability values for individual voxels. One can justify this averaging for the purposes of building a visual model. However, when we consider the actual information content of the model, averaging is an operation that does not preserve the information content in a given set of probability values. As an example, if we consider two voxels, one that has been interrogated a number of times and the other not interrogated at all, we realize that the information contained in these two voxels is very different.

Moreover, simple averaging would tend to give equal weight to each of the values and outliers might swamp the true value. In our formulation, we must be able to weight values according to our measure of confidence in them.

3. The ability to incrementally measure the information content of a collection of voxels.

This requires us to formulate our model such that the information contained in any particular voxel is finite. The backprojection process, corresponding to a particular ping, is an incremental process that should affect only certain voxels. We would like to be able to calculate our global efficiency measure in a manner such that we need only concern ourselves with the immediate voxels

that are affected by a particular ping.

4. Applicability to multiple resolutions.

As pointed out in the introduction to this thesis, sonar sensing encompasses multiple sensors at different levels of spatial resolution. The differences in spatial resolution might arise due to differences in the sensor or from changes in the accuracy of navigation. In either case, the formulation should be able to handle this effect consistently.

5. A measure of the efficiency of a particular sensor in a particular environment.

Sensors are susceptible in different degrees to various effects of the environment. Our measure should be able to evaluate the suitability of a suite of sensors in different environments.

6. Effect of the footprint of the sensor on the mapping measure.

We would like our function to be robust for various kinds of sensors. Specifically we would like it to maintain its characteristics as the footprint along the bottom goes from encompassing single voxels or a very small number, to encompassing a larger area comprising many voxels.

7. The effect of linear operators on the mapping measure.

It is highly desirable to understand the behavior of the mapping measure in terms of linear operators (eg. low-pass filtering) on this measure. This is related to the point we have already made about applicability to multiple resolutions.

3.2 Entropy

Let us now look at the concept of entropy in some detail. The concept of entropy was first formulated by Shannon [98] in an attempt to come to a basic understanding of the processes important to the field of communications. The fundamental issue addressed in his monograph is the

problem of designing a system such that a maximum rate of transfer of information is achieved between a source and a destination. Information is passed between the source and destination by means of symbols across a possibly noisy channel, as depicted in Fig 3.2 below. In general, it is not possible to reconstruct the transmitted signal with certainty by operating on the received signal. Entropy is one measure of the amount of information being produced at the source, being transmitted across the channel and eventually arriving at the receiver.

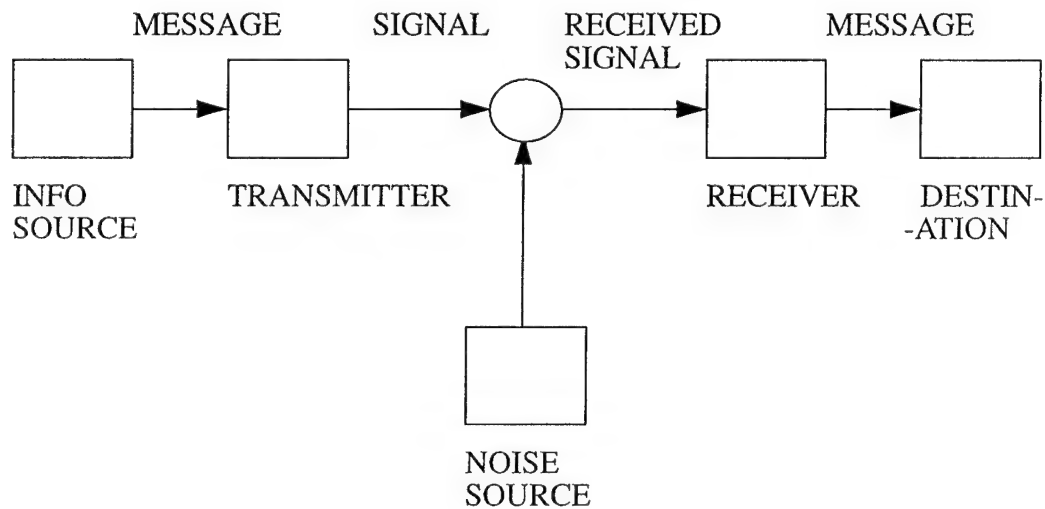


Figure 3.2 A General Communications System

H , the entropy measure for a given set of symbols whose probabilities of occurrence are p_1, p_2, \dots, p_n is given by

$$H = -p_1 \log p_1 - \dots - p_N \log p_N \quad (3.1)$$

Consider the case where two possible symbols, 0 and 1, are being transmitted across a noisy channel at a rate of 1000 symbols per second with probabilities of transmission of each of the symbols,

p_0 and $p_1 = 0.5$. In being transmitted through the noisy channel, 1 in 100 of the symbols is received incorrectly (that is a 1 is seen as a 0 and a 0 is seen as a 1). The rate of information transfer is clearly not the value (990 bits/s) obtained by subtracting the expected number of errors, because the receiver does not know when the errors occurred. To understand this fact a little better, let us carry that kind of reasoning to an extreme by considering a channel in which the received signal is independent of the transmitted signal, and the probability of receiving a 0 is 0.5, irrespective of what was transmitted (and similarly for a 1). Then we cannot say that 500 bits/s were being transmitted, because in reality no information is being transmitted at all, and we could do just as well by flipping a coin at the receiver!

The proper correction to be applied to the amount of information received is a function of the entropy, in this case the conditional entropy, H_y , of the message given our knowledge of the received signal.

In the first example, if a 0 was received, then the *a posteriori* probability that a 0 was transmitted is 0.99 and that a 1 was transmitted is 0.01. The values are reversed if a 1 is received and thus

$$\begin{aligned} H_y &= -[0.99 \log 0.99 + 0.01 \log 0.01] \\ &= 0.081 \text{ bits/symbol or } 81 \text{ bits/s} \end{aligned} \tag{3.2}$$

Thus, we can say that the system is transmitting at a rate of $1000 - 81 = 919$ bits/s.

For our more extreme example the *a posteriori* probabilities are 0.5 for both symbols and

$$\begin{aligned} H_y &= -[0.5 \log 0.5 + 0.5 \log 0.5] \\ &= 1 \text{ bit/symbol or } 1000 \text{ bits/s} \end{aligned} \tag{3.3}$$

So, the rate of transmission becomes $1000 - 1000 = 0$, as we might expect.

3.3 Entropy for Sonars

In considering an entropic measure for sonar sensing the three-dimensional, volumetric probability grids can be considered as discretizations of a given space in terms of our knowledge regarding the presence or absence of material, with the probability value at each discretized voxel a measure of the confidence of our knowledge at that particular voxel. Consider the sonar transducer to be a source of information about the environment. Note that various factors, such as navigation and attitude errors, act on the individual probability values for each of the voxels and corrupt the true value. We would like specifically to be able to quantify the validity of our data in the face of these errors, enabling us to obtain a good measure of the efficiency of our mapping process.

Besides Stewart [105], Saridis [96] also examined the concept of entropy for robotics. He proposed organizing an intelligent robot in terms of a three-level hierarchical structure representing its Organization, Coordination, and Execution levels and using entropy as a common measure of the uncertainty at each of these three levels. He cast the problem of intelligent control as “the mathematical programming solution that minimizes the total entropy.” His entire focus has been the control-design problem and he assumes the existence of a database of events with associated probabilities. Moreover, he does not address the issue of going from the raw sensor data to a probability model. Moravec [82] also defines the concept of entropy for a probability grid but does not elaborate on its utility.

Let us look closely at how we might map the sonar sensing modality onto the information-theoretic concept of entropy. Returning to the general communications system described earlier, consider a mapping such that the environment can be considered to be our information source. The sonar sensing modality and the sources of noise inherent in the sensing process can be considered

in combination to be the transmitter, the noisy channel, and the receiver. Finally a particular ping and the probability values that we assign to our sensor cone may be thought of as the message arriving at our model of the environment, which we identify with the communications-system equivalent of the destination.

Entropy is a measure of the information content of our model of the environment. Thus, as the information from more pings is added to our model, its entropy should decrease. Eventually the entropy should converge to a steady state where it approaches a minimum for the particular set of operating characteristics. Figure 3.2 plots out the entropy for a single voxel based on our concept of occupancy and emptiness. We see that as our information regarding the voxel goes up, that is we become more confident about the occupancy or emptiness for that voxel, its entropy decreases.




Figure 3.3 Entropy for a Bi-valued Distribution

Here it may be pointed out that we have made an implicit assumption regarding the ergodic

nature of the environment, when mapped with the sonar sensing modality, to allow us to place the whole problem in an entropic framework. Cobra [30] modelled bottom backscattering much more restrictively as a wide-sense stationary random process in successfully deriving an estimate of geometric distortions in sidescan sonar imagery. If we consider the sonar sensing modality in an unstructured underwater terrain the process is indeed ergodic. Moreover, in the case of a terrain where there are significant changes in the morphology, we can divide the area up into regions that are locally ergodic.

3.4 Suitability of an Entropic Efficiency Measure

Let us run down some of the desired characteristics that we outlined earlier and thus examine the suitability of an entropic efficiency measure.

3.4.1 Incrementally measuring information content

The information content of a volumetric collection of voxels changes with each ping such that the change in entropy is caused only by the voxels ensonified by that particular ping. Thus, as desired, our value of the global entropy is only a function of the local voxels that were affected by a particular ping.

3.4.2 Multisensor Sonar Modelling

A typical underwater mapping scenario might map out the area with a ship-borne sensor, as a first step, then follow it up with a finer scale survey using a towed or remotely operated vehicle. The towed or remotely operated vehicle might contain one or more sensors, which might be used

depending upon the requirements for the survey. We would like to utilize the overlapping data in a consistent framework. One of several ways of doing multisensor modelling is to do a band-limited interpolation from a lower resolution to a higher resolution over the volume or area being mapped. Rather than look at this problem specifically, we focus instead in the next section on the much more general problem of applying linear operators to our entropic model.

3.4.3 The Effect of Linear Operators

Besides the applications for multisensor modelling pointed out above, linear operators often find use in typical AUV applications. For example, in path planning one might need to perform smoothing operations, a form of low-pass filtering, on a raw map. Thus, it would be beneficial to understand the relation between the entropy of the map before and after the application of a linear operator.

So far we have only concerned ourselves with a discrete set of probabilities, $p_1 \dots p_n$. For the purpose of a better insight into entropy and its relation to filtering operations, let us now define the entropy for a continuous distribution with probability density function given by $p(x)$ as:

$$H = - \int_{-\infty}^{\infty} p(x) (\log p(x)) dx \quad (3.4)$$

where x may be multidimensional.

While the properties of the discrete and continuous entropies are similar, one important difference is that in the discrete case the entropy is an absolute measure of the randomness of the random variable; in the continuous case the measurement is relative to the coordinate system.

If we change coordinates, in general the entropy will change. If we go from $x = (x_1 \dots x_n)$ to $y =$

$(y_1 \dots y_n)$ the new entropy is given by

$$H(y) = - \int_{-\infty}^{\infty} p(x) J\left(\frac{x}{y}\right) \left(\log \left(J\left(\frac{x}{y}\right) p(x) \right) \right) dy, \quad (3.5)$$

where $J\left(\frac{x}{y}\right)$ is the Jacobian of the coordinate transformation. We can manipulate the previous equation to show

$$H(y) = H(x) - \int_{-\infty}^{\infty} p(x) \left(\log J\left(\frac{x}{y}\right) \right) dx \quad (3.6)$$

That is, the new entropy is the old entropy less the expected logarithm of the Jacobian. In the continuous case the entropy can be considered a measure of randomness relative to an assumed standard, namely the coordinate system chosen with each small volume element $dx_1 \dots dx_n$ given equal weight. On transforming $x \rightarrow y$ we are now measuring the entropy in a coordinate system where each of the volume elements $dy_1 \dots dy_n$ are given equal weight.

A particular case of changing coordinates is the linear transformation

$$y_j = \sum_i a_{ij} \cdot x_i \quad (3.7)$$

In this case the Jacobian is simply the determinant $|a_{ij}|^{-1}$ and

$$H(y) = H(x) + \log |a_{ij}| \quad (3.8)$$

For measure-preserving transformations such as rotation, $J = 1$ and $H(y) = H(x)$.

As another example we consider the entropy of our sensed volume before and after we filter it with a filter having spatial frequency response $Y(f)$. The operation of the filter can also be thought of as a linear transformation of coordinates. If we consider the original volume and its various frequency components as making up the original coordinate system, then the new frequency compo-

nents are simply the old frequency components multiplied by factors given by the filter. The coordinate transformation matrix is thus essentially diagonalized in terms of these coordinates. The Jacobian of the transformation is (for the n Fourier components)

$$J = \prod_{i=1}^n |Y(f_i)|^2 = \exp \sum \log |Y(f_i)|^2$$

where the f_i are equally spaced through the band W .

In the limit the expression above becomes

$$J = \exp \frac{1}{W} \int_W \log |Y(f_i)|^2 df \quad (3.9)$$

which leads to the result

$$H(y) = H(x) + \frac{1}{W} \int_W \log |Y(f_i)|^2 df \quad (3.10)$$

Note that in both cases above the change in entropy associated with the transformation may be precomputed and is fixed for a particular transformation or filtering operation.

What these results seem to suggest is that entropy is a function of the coordinate system and that our measure of efficiency of the mapping process is as a result, rather inconveniently, a function of our choice of coordinate system. However the crucial point to note is that the information rate of our sonar sensor depends upon the difference in entropies, and this difference does not depend upon the coordinate system, as each of the coordinate systems are changed by the same amount. Thus, from the standpoint of AUV applications, namely measures of efficiency of mapping, our sensors are independent of changes in coordinate systems due to the application of linear operators.

3.4.4 A Measure of Sensor Efficiency in a particular Environment

Given this formulation of Entropy for underwater mapping, let us now apply it to some of the examples mentioned earlier. Note that entropy can be applied on a per-ping basis or on the global volume. At this juncture we turn our attention to examining the effects of an individual ping on uncertainties in navigation and attitude.

We form our sensor cones in the manner described in Section 2.4 for a pencil beam sonar of beamwidth 1.7° (i.e. with the characteristics of the Imagenex pencil-beam sonar) then calculate the entropy for each ping. We plot the change in entropy per ping versus the error for the cases with navigational and attitudinal uncertainties.

Figure 3.4 shows the effect of navigational uncertainties in a manner consistent with our intuition of this concept. The information per ping falls off smoothly as the navigation error increases.

Figure 3.5 shows the effect of attitudinal uncertainty. An obvious feature in this graph is the sudden discontinuity between the original and the error entropies. By examining the original graph one might suspect that this is caused by an aliasing effect caused by sampling a 1.7° beam close to its source.

Figure 3.6 shows the effect of the attitudinal uncertainty by plotting the entropy of occupancy versus attitude error. At a range of 75 m, i.e. far from the source, the effect of the aliasing due to sampling is not as evident as in Fig 3.4. We can also compare this graph with the qualitative graph for attitude uncertainties obtained by Stewart [106] and see that the basic shape is in close agreement.




Figure 3.4 Entropically Measuring Navigational Uncertainty




Figure 3.5 Entropically Measuring Attitudinal Uncertainty




Figure 3.6 Entropic Attitudinal Error for Occupancy

3.4.5 Consistently Assimilating Conflicting Information

A basic assumption in probabilistic volume modelling is that regions of the volume under consideration will be ensonified multiple times. In reality, this is usually the case, as a particular voxel is often interrogated several times. The issue then becomes how to incorporate these multiple, possibly inconsistent values for a particular voxel.

There are many possible sources of inconsistencies. A reality of imaging underwater that we have not discussed so far is that our sensors cannot resolve temporally dynamic events because of fundamental sampling constraints. Unfortunately, in the pursuit of our mapping goals, we often

come across situations where highly dynamic events do occur, a good example being ensonification of a fish by the sonar. This has the consequence of biasing our representations of the world with data that might be viewed as inconsistent.

Other inconsistencies may also arise because of differences in the proximity of the sensor whose beam we are modelling. As we move closer, we can localize a target much better in a fixed grid because the beam has not spread out as much as it would over a greater distance.

These inconsistencies show up in our mapping efforts and we would like our entropic measure of efficiency to be robust to their effects. Given the huge dynamic ranges of the interacting quantities in the sonar equation, we would also like to take advantage of those sensor readings that are less susceptible to noise. Our rule of combining probability values should have the following characteristics.

1. It should be able to combine multiple values for the same voxel such that repeated evidence in favor of occupancy or emptiness leads to stronger belief in each of these respective propositions.
2. The rule should weight values closer to 0.0 and 1.0 much more than the values near 0.5, with the value at 0.5 being treated as an identity operator.
3. The rule should be able to arbitrate between inconsistent values that point to occupancy and emptiness for the same voxel.

We use a simple form of the Dempster-Shafer rule of combination [38][97][128]. Our algorithm has two distinct cases that present themselves: 1) When we are provided with new evidence for a particular voxel which supports the evidence for occupancy or emptiness already associated with that voxel; and 2) when the new evidence is in direct conflict with the evidence in a particular voxel.

The first case has been studied extensively. Bernoulli [10] looked at the problem and his work was extended by Lambert [70] in the so called Bernoulli's Rule, which we outline below.

If p_1 is the probability supporting a particular hypothesis and p_2 another entirely separate probability supporting the same hypothesis, then the two probabilities support the hypothesis in combination with a probability p , given by

$$p = 1 - (1 - p_1) (1 - p_2) \quad (3.11)$$

We note that our formulation for stochastic modelling is not the same as the formulation for belief required for use with the Dempster-Shafer Rule. The Dempster-Shafer rule supposes the probabilities corresponding to our belief are in the range 0.0 to 1.0, where 0.0 implies we have no evidence to support the particular belief and 1.0 implies we are absolutely certain in that particular belief. Our stochastic framework is different and assumes values of 0.5 represent no information and goes in either direction towards 0.0 and 1.0 depending upon our belief regarding emptiness or occupancy. But it is easy to scale our framework to apply the Dempster-Shafer combination rule. Thus, for the case where the probabilities both support emptiness or occupancy, (3.11) is transformed to Eqns (3.12) and (3.13) respectively

$$p = 2 p_1 p_2 \quad (3.12)$$

$$p = 1 - 2 (1 - p_1) (1 - p_2) \quad (3.13)$$

In Fig 3.7 we plot out the results of combining a series of values, with probabilities ranging from 0.0 to 1.0 in steps of 0.1, and the manner in which they support our beliefs about occupancy and emptiness. We see that the curves all approach complete confidence in their respective beliefs at rates in accordance with the incoming values.

What remains is for us to reconcile conflicting evidence supporting both occupancy and empti-

ness for a particular voxel. We expect conflicts to occur for reasons pointed out earlier in this chapter. The Dempster-Shafer rule treats conflicting evidence in support of two mutually exclusive events by lowering appropriately the evidence supporting each of the events. We could use this approach by keeping separate tabs on the evidence for occupancy and emptiness. But at some stage we would have to resolve conflicts, if any, between the evidence supporting occupancy versus the evidence for emptiness in order to produce a final map. So, rather than wait, we simply make the choice to resolve conflicts as they occur by linearly averaging out the two values (3.14).

$$p = (p_1 + p_2) / 2 \quad (3.14)$$




Figure 3.7 The Dempster-Shafer Rule of Combination modified for Stochastic Modelling

3.5 Entropies of Emptiness and Occupancy

A truth pointed out long ago by the Artificial Intelligence community is that it is just as important

to know that an environment is free of obstacles as it is to know the location of obstacles. Knowing that a particular area is devoid of obstacles is very different from not knowing anything about the area under consideration. This fact is especially important in AUV path planning where the safe functioning of the vehicle over-rides all other concerns.

Our intuition about mapping, however, is based on a slightly different set of priorities. While we would appreciate information regarding the absence of material, our intuition has been honed on years of making maps of features that exist.

The measure of entropy measures both the entropy of occupancy as well as the entropy of emptiness. Note that the entropy of a ping for a particular sensor is usually much greater for ranges approaching the limits of the sensor and that a blind use of entropy for global level mission planning as it has been formulated so far might easily lead to an optimum mapping strategy of staying in the water column!

Clearly what is called for is keeping track of the entropies of occupancy and emptiness independently with the mapping priorities more heavily slewed towards looking at the entropy of occupancy and overall vehicle path planning utilizing the two in some appropriate combination which weights the requirements for getting good data with issues of vehicle safety.

3.6 An Independent Validation: The 10% Rule

Whenever one lays out an analytical scenario for some aspect of reality an issue is always whether one has modelled all relevant aspects of a phenomenon. We will find out, sooner or later, as our concept of reality goes awry because of an invalid assumption or, in a much happier circumstance, our model of reality works elegantly on its own and has applicability in ways we did not originally intend in our formulation.

Among side-scan sonar operators an empirical rule of operation has always been that a tow fish should be flown at an altitude which is 10% of the swath width to achieve the best imaging performance. As a test to our formulation of entropy we decided to evaluate this empirical rule of thumb.

We assume a flat bottom and a target strength corresponding roughly to that of mud. For a constant swath width we vary the altitude of the towfish and plot it against the entropy of occupancy, calculated for each altitude. The results, shown in Fig 3.8, show an optimum altitude of roughly 7 percent. The existence of an optimum close to that predicted by the empirical rule of operation serves as an interesting independent validation of our modelling formulation.

We can offer an explanation after the fact. The value of imagery is a function, over the entire swath width, of the backscattered energy as well as the actual energy arriving at that particular point. As we increase the altitude at which we fly, the grazing angle becomes less acute, allowing for a much greater degree of backscatter. At the same time, the transmission losses increase as sound now must travel much further through the medium before ensonifying the bottom. The shape of the typical backscatter versus grazing angle curves determines the optimum height.




Figure 3.8 Entropy of Occupancy as a function of Altitude for a Given Swath Width

CHAPTER 4

Modelling Sensor Deployment

We now look at using the concepts developed in the previous two chapters to understand some of the issues involved in deploying these sensors on real AUVs. We have already noted that power is one of the most constraining factor for AUVs today. Bradley [18] derives an interesting relation (reproduced in Appendix A) between the energy needed for vehicle propulsion, the energy corresponding to the hotel load, and the ultimate range of the vehicle. He points out that for every hotel load there is an optimum speed that gives a maximum range as shown in Fig 4.1. The geometric parameters used in this figure are for a NACA series 58 body of revolution with a length to diameter ratio of 5:1. The co-efficient of drag is based on actual measurements [57] on the Fast Hydrographic Profiler [19].

One of the implicit assumptions in his analysis is that the rate of information flow from the sensor on the vehicle is constant for the speeds under consideration. This assumption may or may not be true for different types of sensors. It is definitely not true for high frequency acoustic sensors. We look to an entropic framework to answer questions regarding the optimum speed we must travel to maximize the rate of information for a particular mission. We first examine the rate of change of entropy for repeated pings from the same position in order to quantify the redundancy required for optimal mapping. We then use this to understand how our information flow is affected by changing the speed of our AUV. Finally, we look at how navigation uncertainties manifest themselves in choosing the optimum speed.




Figure 4.1 Range vs Speed for Different Hotel Loads

4.1 Entropy Change for Repeated Pings

One of the basic assumptions in using the stochastic modelling methodology is that the area of interest will be ensonified multiple times. In the previous chapter we looked at combining multiple probability values for a single voxel. We now look at the change in entropy on a per ping basis.

For this purpose we combined multiple pings, from a pencil-beam and an broad beam sonar, at a representative range (chosen to be 20m — dictated for most AUVs by the need to do simultaneous optical sensing). The results, graphed in Fig 4.2, show the change in entropy for the two sensors as a function of ping number. We observe that a small amount of redundancy (3-4 pings) in pinging the same area greatly increases the information on a per-ping basis. If we continue to ping further the amount of information per ping rapidly diminishes. The differences between the two curves are dominated by the differences in the source levels and differences in the beam patterns.




Figure 4.2 Entropy Change for Repeated Pings

4.2 Effect of Speed on Sensor Efficiency

In the light of the discussion in the previous section we now turn our efforts towards understanding the optimal speed for a particular sensor mounted on an AUV. To do so we simulated a realistic bottom profile and flew above it using a path planning algorithm based on recursively filtering the altitude obtained from our sensors as shown in Fig 4.3.




Figure 4.3 Simulating Different Speeds for Optimizing Sensor Efficiency

The change of entropy for various speeds was normalized with respect to time to obtain the entropic change per unit time. In plotting Figs 4.4 and 4.5 we further normalized the results, linearly about the maximum entropic change for each sensor, to obtain a measure for the efficiency in mapping. The overall trend in these curves is opposite to the trend in the range versus speed curves plotted in Fig 4.1. Figs 4.4 and 4.5 point towards a faster optimal speed for the 10W echo sounder relative to the 25W pencil-beam sonar. It turns out that the pencil-beam sonar is very close to the optimum speed for both the curves but the overall curve is in opposition. We discuss the significance of these results at the end of the chapter in section 4.4.




Figure 4.4 Normalized Entropy Change vs Speed: Pencil-Beam Sonar




Figure 4.5 Normalized Entropy Change vs Speed: Broad-Beam Echo Sounder

4.3 The Effect of Navigation Uncertainty

We can follow the procedure in the previous section to incorporate navigation uncertainty into our calculations. By modelling the beams, as described in section 3.4, we can rerun our simulations to obtain the results in Figs 4.6 and 4.7 for a navigation uncertainty of 5m (typical of a low frequency long base line navigation system). As one would expect the navigation uncertainty tends to favor higher redundancy in our mapping efforts which in turn lowers the optimal speed at which we must travel. Further, we see that the echo sounder, because of its broad beam, is less susceptible to the navigation uncertainty, and thus the change in speed corresponding to that sensor is lower than the relative change in speed for the pencil beam sonar.




Figure 4.6 Normalized Entropy Change vs Speed:
Pencil-Beam Sonar w/ Navigation Error




Figure 4.7 Normalized Entropy Change vs Speed:
Broad Beam Echo Sounder w/ Navigation Error

4.4 Optimizing the Overall Efficiency of Mapping

The difference in the results between Fig 4.1 and the Figs 4.4 through 4.7 arises from the assumptions on which each is constructed. Fig 4.1 optimizes the propulsion load versus the hotel load. It partitions the energy equally between propulsion and sensing so that a higher hotel load leads to higher propulsion loads and correspondingly higher optimal speeds. Figs 4.4 through 4.7 on the other hand are based on our models of the sensors.

Note that from a global mapping standpoint if we are moving slower than the optimum speed we can choose to subsample our data, or if the hardware allows to ping only at the effective optimal rate. However, if we are moving much faster than the optimal speed the solution is not as easy. The separation between successive footprints of the sensor along the bottom becomes increasingly non-

redundant and eventually the sampling along the bottom becomes too coarse for the grid onto which we are backprojecting our sensor cone.

Note that for the pencil-beam sonar the normal mode of operation calls for the head to be mechanically scanned from side-to-side as the vehicle moves along the bottom. In this case given the 0.5-2 m/s operating speeds of AUVs and ROVs we are definitely in a non-redundant sampling regime where several grid cells are not being sampled at all. This may be acceptable for map-making (in the manner described in Appendix C) but, according to our stochastic formulation, will not have sufficient redundancy for change detection applications.

The discussion in the previous paragraph points out the utility of sidescan sonar systems. A side-scan sonar may be modelled as a collection of fixed pencil-beam sonars. The wide beam in the cross-track direction means that at normal operating speeds we can achieve the three or four fold redundancy that is optimum for our mapping efforts. The increase in swath coverage (three orders of magnitude) more than compensates for the increase in the hotel load. This is evident from (4.1), derived in Appendix A, which shows a very weak dependence of speed on P_H . The optimum speed goes up only as 1/3 power of P_H .

$$u = \left(\frac{P_H}{2C_d S_0 \rho L^2} \right)^{1/3} \quad (4.1)$$

The 75-100W hotel load corresponding to sidescan sonars is possible with small AUVs today. The stumbling block behind the immediate deployment of such sensors on small AUVs is the size and the weight of the sensor and its associated electronics. For the larger AUVs currently being planned or in operation some form of sidescan sonar is undoubtedly the sensor of choice.

CHAPTER 5

Application Data Sets

We now look at applying our methodology to real data sets. The first set of data was obtained during mapping exercises off the Woods Hole Oceanographic Institution dock. A variety of high-frequency sensors were mounted on the JASON ROV to examine the issues behind high frequency, high-resolution mapping in a shallow water environment. The second set of data was obtained during mapping exercises using a broad-beam, down-looking, echo sounder mounted on the ABE AUV.

These two datasets are interesting as they represent unique and different issues in doing high frequency mapping with underwater vehicles. Working with these two sets of data allows us to examine the basic limitations and algorithms for high frequency acoustic mapping as well as some of the important issues for deployment on real AUVs.

5.1 Shallow-water mapping

There were several reasons for choosing to focus on high-frequency acoustic mapping in a shallow-water environment. While high-frequency mapping from underwater vehicles is a challenging task in itself, it is even more challenging in a shallow-water environment. Recent military conflicts have also underscored the need for fine-scale mapping in such an environment.

Another major consideration in using this dataset is that a shallow-water environment allows us to ground-truth our mapping efforts. We could measure the locations of various features in our imagery, such as dock pilings, precisely and physically were able to walk out and independently

measure these very structures. This is in sharp contrast to work in the deep ocean where high-frequency mapping exercises can be carried out, but the supporting ground truth information to help characterise the success of our mapping efforts can only be qualitative in its nature.

This work also has significance in that we can return to the area after considerable periods of time and repeat our mapping efforts to help understand spatiotemporal variability. This is important not only from a military viewpoint but is also part of the basic mission of AUVs. A good understanding of spatiotemporal variability will be of considerable help in actual long-term deployments at hydrothermal vents and other areas of interest.

5.1.1 Imagenex Sensor Characteristics

The sensor used in our mapping exercises with JASON is an Imagenex 675 -kHz pencil-beam, mechanically scanned sonar. The transducer is circular with a diameter of 3.5 inches forming a conical (pencil beam) beam pattern of 1.7° angular extent. The beam we modelled and analyzed in Section 3.4 is based on the characteristics of this sensor.

The sensor was mounted on the front of the vehicle on the port side so the axis of the beam was parallel to the horizontal plane. We settled on using this look-ahead mode so we could image objects repeatedly, such as the pilings of the dock and a nearby pier. By flying the vehicle at different altitudes we could effectively obtain slices of the 3-dimensional volume that we set out to map.

5.1.2 Navigation and Attitude

Figure 5.1 is a plot of the navigation tracks. In plotting the vehicle navigation and attitude we attach markers A through F to identify points that coincide with changes in the vehicle tracks.

These tracks were laid out to parallel the geometry of the area as far as possible given the minimum depth required for vehicle operation. We plot the positions in world coordinates for each of the two transducers mounted on the vehicle with dotted lines. The two SHARPS navigation transducers were mounted close to the center line, fore and aft of the vehicle. The transducer on the stern was very close to our arbitrarily chosen coordinate centre of the vehicle. The physical separation between the two dotted lines in the figure is caused by changes in the vehicle's heading.

The vehicle was always pointed towards the features of interest. In moving from A to B the vehicle was pointed perpendicular to the dock and moving forward, so the fixes plot out in space close to one another. In moving from position B to position C the vehicle was still perpendicular to the dock and but moving sideways so that the the fixes are separated from one another.

These vehicle fixes were transformed to a common origin on the vehicle and merged with a Kalman filter that merged in the velocities obtained from an Acoustic Doppler Current Profiler, to form a composite estimate of vehicle position, plotted in the figure as a dashed line.

Figure 5.2 shows the position in Z as a function of time. We see that the vehicle was having difficulties maintaining a constant depth. The oscillatory behavior in Z was a function of the fact that new, more powerful thrusters had just been installed on the ROV but the software for the control algorithm to do the depth servoing had not been updated.




Figure 5.1 JASON Navigation for the Shallow Water Survey operation




Figure 5.2 JASON Depth for Shallow Water Mapping

Figure 5.3 shows representative JASON pitch, roll, and heading plotted as functions of time. We see that the roll and pitch were of small amplitude and seem to be well characterized by the sensors along the respective axes. The offsets on the pitch and roll can be attributed to the presence of a number of other sensors and a robotic arm that were installed on the vehicle and tended to weigh down the stern, starboard side of the vehicle. The heading servo is seen to be extremely well-behaved and our measurements of heading seem to characterize vehicle behavior along this axis exceedingly well.




Figure 5.3 JASON Attitude (degrees) vs Time (samples at 2 Hz)

5.2 Computational issues in 3D mapping

Let us now focus on some of the computational issues in 3-dimensional mapping. From an AUV standpoint, algorithms for mapping must be realizable in real time within the limited computational resources on board the vehicle. Real-time mapping allows us to deal with large amounts of data that characterize the sonar sensing modality. It allows us to compress the bandwidth of our sensors so that it becomes feasible for us to transmit the data to the outside world through a low-bandwidth acoustic modem or simply to record the data in a compacted form.

As long as we are capable of processing the pings in real time, the memory requirements for the volume being imaged are bounded by the resolution of the volume being imaged. In this respect we can be within the envelope of the computational resources available. Given that we can bound the memory requirements, the other concern is computational complexity. A 3-dimensional back-projection algorithm implies the ability to backproject a sensor cone from an arbitrary position in the volume along an arbitrary direction. The computations involved in achieving such a goal are significant. The path we chose to follow entails precomputing key computationally expensive operations. Thus, instead of computing on the fly we accessed a look-up table and read out the required values — in essence trading computational complexity for memory requirements.

Specifically, one of the major computations is in identifying the voxels that are ensonified by a particular ping. This typically involves computing several transcendental functions for a large neighbourhood and is one of the major bottlenecks in achieving real-time performance. So, we precompute all the voxels that lie within broad, overlapping cones for the complete range of sensor directions as illustrated in Fig 5.4. The overlap ensures that, for a particular angle, the required sensor cone is completely encompassed by a single precomputed broad cone.




Figure 5.4 Precomputed Overlapping Sensor Cones

We settled on precomputing sensor cones of 5° width at 2.5° intervals. We needed only to compute the sensor cones, corresponding to the in-plane θ direction, from $0-87.5^\circ$. The sensor cones for the remaining three quadrants can be derived from these by changing the signs of the values in x and y depending upon the quadrant in which the sensing cone lies.




Figure 5.5 Variation in phi for the Sonar Sensor Beam

To accomplish full 3-D mapping we utilized overlapping sensor cones for θ and ϕ . We note that while the mechanically scanned sensor had complete freedom with respect to rotation about θ , the variations in ϕ were due to vehicle attitude changes in roll and pitch.

By examining the dynamics of the vehicle one can estimate the range of motion in ϕ . In Fig 5.5 we plot the phi dependence of the sensor cone for our dataset. Thus we chose to compute ϕ over a range 0-12.5°.

With a sensor cone encompassing 5° in each of the phi and theta dimensions, for a typical volume comprising 400 x 400 x 96 voxels to achieve a 25-cm resolution in the x and y dimensions and a 12.5-cm resolution in z, we wind up with over 30,000 voxels per sensor cone. This is a large number of voxels to handle per ping. To further simplify our processing we note that our range returns are distributed from 0-65 m but that the majority of the range returns come from a range less than 20 m as can be seen in the histogram in Fig 5.6. We thus chose to sort the precomputed voxel values by range. This has the effect of substantially cutting down the voxels we examine, and the overall performance of our algorithm now meets and exceeds the criterion for real-time mapping.




Figure 5.6 Histogram of Returned Ranges

5.3 Mapping Results

In this section we present some of the results of using our stochastic modelling framework for 3-dimensional mapping. For reference, Fig 5.7 shows an aerial view of a portion of the Woods Hole Oceanographic Institution dock which was the subject of the first of our mapping exercises.

In the foreground one can see the L-shaped, so-called MBL pier, the seawall and the dock itself. The dock is a concrete structure supported by pilings. It is designed to serve large ocean-going vessels, and the depth varies from about 10 cm at low tide near the seawall to 15 m along the dock as we move out towards open water.

Our mapping efforts in this area have been spread over a couple of years and have used a variety of sensors. In the fall of 1993 a 200-kHz sidescan sonar was attached to a calibrated stepping motor

which was stepped in a circular pattern to obtain the image in Fig 5.8. This was an interesting exercise as it provided us with an image which was devoid of the typical effects of platform attitudinal motion as well as the uncertainties due to navigation. The MBL pier, the sea wall, and the pilings under the dock can be clearly seen.

In the fall of 1994 we returned to the site in efforts to characterize the effects of temporal variability as well as help us quantify the basic sensor limitations when used in real operating circumstances on a real ROV.

The results of our 3-dimensional stochastic modelling efforts with the 675-kHz pencil-beam sonar are shown in Fig 5.9. All occupied voxels were overlaid on the 200-kHz sidescan sonar data which was texture-mapped on the Seabat 455-kHz multibeam profile data. Some other close-up views of our mapping efforts are shown and discussed in the next section. For comparison the results of geometrically correcting the same data are shown in Appendix C.

In this image we would like to bring attention to the rapidly shallowing segment in the corner contained by the MBL dock and the seawall and the very good qualitative agreement for the three different sets of data. We discuss the significance of these results for change detection in the next section.




Figure 5.7 Aerial View of the WHOI dock




Figure 5.8 200-kHz Side-Scan Sonar Data (Mechanically Scanned)



Figure 5.9 3-D 675-kHz Imagenex Data

5.4 Entropy as a measure of stability and convergence

We now look at the mapping results discussed in the previous section from an entropic standpoint. A basic issue we wished to explore in this thesis is the convergence and stability of the stochastic modelling framework and the applicability of our framework for change detection.

In Fig 5.10 we plot the global entropy versus time for all voxels in our volume. The entropy is seen to decrease monotonically. The reason for this monotonic decrease can be seen quite clearly in Fig 5.11 where we show some of the slices in our volume. The speed at which we are travelling in combination with the mechanically scanned sensor implies that we are in the regime of non-redundant coarse sampling with respect to the 3-dimensional volume. Further, the offsets in roll and pitch imply that our successive sensing cones are not confined to a particular plane tend to scatter a particular sensing cone across multiple planes leading to the sparsely ensounded slice in Fig 5.11.




Figure 5.10 Change of Entropy vs Time

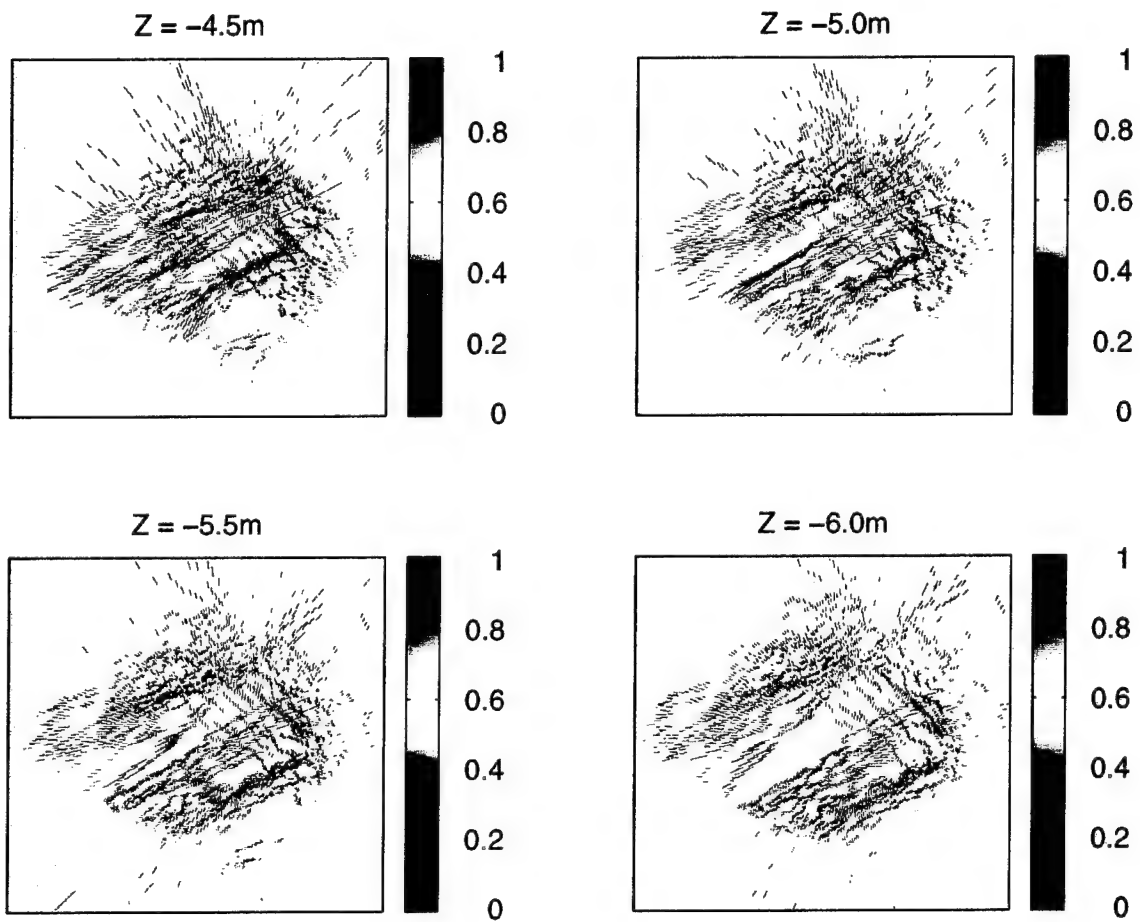


Figure 5.11 Volumetric Slices for the Stochastic Volume

The results are quite different when we focus on a subsection of our data flown at roughly constant depth. The depth for part of this data is plotted in Fig 5.12. We do see the oscillations in the depth because of the control software but the mean value is roughly constant. We utilized a “pseudo-2-dimensional approximation” of the true volume. A simple 2-dimensional approximation, that is, assuming that successive sensing cones were limited to a plane, tends to smear out the features of interest due to vehicle roll and pitch. Our pseudo-2-dimensional approximation instead, calculates the entire sensing cone in 3D but projects the results onto a 2-dimensional plane. Now when we plot out the global entropy for our 2-dimensional slice (Fig 5.13) the entropy is seen to be converging. This can also be seen in fig 5.14 where we plot out the entropy change on a per-ping basis. We notice that as we converge individual pings may backproject conflicting data and an overall positive change in entropy. We hypothesize that this effect is at least in part due to the oscillations in the depth. Overall though, the entropy is converging.




Figure 5.12 Depth for Pseudo 2D Mapping




Figure 5.13 Change of Entropy for Pseudo 2D Mapping




Figure 5.14 Change of Entropy vs Time

In Figs 5.15 through 5.18 we plot out the occupied voxels as we conduct our pseudo-2-dimensional mapping exercise. These figures show that we are indeed collecting real information as we continue to ping. They also refute the criticism that probabilistic mapping tends to lowpass filter the results of mapping to the point that feature-based approaches cannot be used in combination with probabilistic mapping. As can be seen in this set of figures, the MBL pier and pilings underneath the dock can be clearly distinguished.




Figure 5.15 Pseudo 2D Mapping - 400 pings




Figure 5.16 Pseudo 2D Mapping - 800 pings




Figure 5.17 Pseudo 2D Mapping - 1200 pings




Figure 5.18 Pseudo 2D Mapping - 1600 pings

5.5 Acoustic Mapping with ABE

We have examined the concept of entropy for 3-dimensional acoustic mapping in some detail. Let us now examine how we might use the concept of entropy for real AUV operations. Given our entropic insight into acoustic mapping, we are now in a position to immediate sensing strategies in terms of the actual information flow from our sensor.

Given that most AUVs are still in a developmental stage and that very little effort has been expended in carrying out acoustic mapping with real AUVs, our analysis is limited but we lay out a framework for further work.

The ABE AUV is a good platform for acoustic imaging purposes. It has been primarily designed for conducting deep-ocean imaging surveys. Its designers realised that one of the prime consider-

ations for an imaging platform is stability. Thus, they chose a geometry that has wide separation between the center of buoyancy and the center of gravity and this wide separation means that the vehicle is passively stabilized in pitch and roll.

ABE also possesses a good power management system that allows us to accurately monitor the energy consumption by the vehicle. The power management system operates at a relatively high bandwidth and in combination with the activity log of the vehicle allows us to accurately profile the power consumption of each of the sensors as well as different modes of propulsion.

We carried out a simple acoustic mapping experiment by mounting an echo sounder on ABE, of the type typically used on sailboats. The particular transducer is a broad-beam design, operating at 200kHz, and manufactured by Robertson Corporation. The transducer was chosen because it is cheap, has a simple interface, and modest power requirements. Our acoustic mapping work was carried out at a neighbouring lake. The geometric results obtained in our mapping exercise are shown in Fig 5.19.

We also modelled the echo-sounding transducer stochastically (this was the broad beam transducer modelled in Section 2.4) and thus could measure the change of entropy over time during our mapping exercises (plotted in Fig 5.20). We notice that the change of entropy is dependent on two factors. The first is simply the altitude at which the vehicle is flying and the second is the redundancy in our data. The value of energy expended during the entire mapping exercise was logged by the energy consumption profiling module on ABE, and the results are plotted in Fig 5.21.

By simply dividing out the values in these two graphs we obtain an instantaneous rate of energy consumption per bit of acoustic mapping information, as shown in Fig 5.22. The top plot shows large energy-per-bit spikes, which correspond to very little changes in entropy when we obtained redundant data. The bottom plot is a zoomed in section of the top plot and shows the dependence

of energy per bit on altitude.

As we stated earlier these results are crude, at best, and considerable work still needs to be done. However these simple results do show the utility and the possibilities of applying our entropic framework to the task of acoustic mapping with AUVs.




Figure 5.19 Geometrically Derived Depth Map for ABE at Herring Pond




Figure 5.20 Change of Entropy for Depth Mapping at Herring Pond




Figure 5.21 Energy versus Time for Depth Mapping




Figure 5.22 Instantaneous Energy per bit

CHAPTER 6

Post Mission Analysis

So far we have looked at the applicability of an entropic framework to various tasks associated with active sensing from an AUV. In Chapter 4 we looked at using the concept of entropy for optimizing the speed of the AUV. In chapters 3 and 5 we looked at using entropy for resolving inconsistent information, for examining the stability and convergence of our mapping methodology, and for local sensor optimization. In this chapter we examine how we might use the stable, convergent map we have constructed for change detection and post-mission analysis such as path planning for future missions.

6.1 The Analysis of a Convergent Map

Figure 6.1 summarizes the use of entropy within a stochastic modelling framework. Our transducer interacts with the environment as we form a sensor cone for a particular ping. This sensor cone takes into consideration the inherent errors in navigation and attitude that may exist for the particular set of environmental characteristics. A measure of the entropy of the sensor cone at this juncture characterizes the efficiency of our sensor in the operating environment. We then back-project the sensor cone into the map by combining it with the available navigation and attitude information and calculate the entropy of the map as a whole. In chapter 5 we showed that our mapping efforts yield a model that is stable and converging. We now take a closer look at the entropy of the stable convergent map.

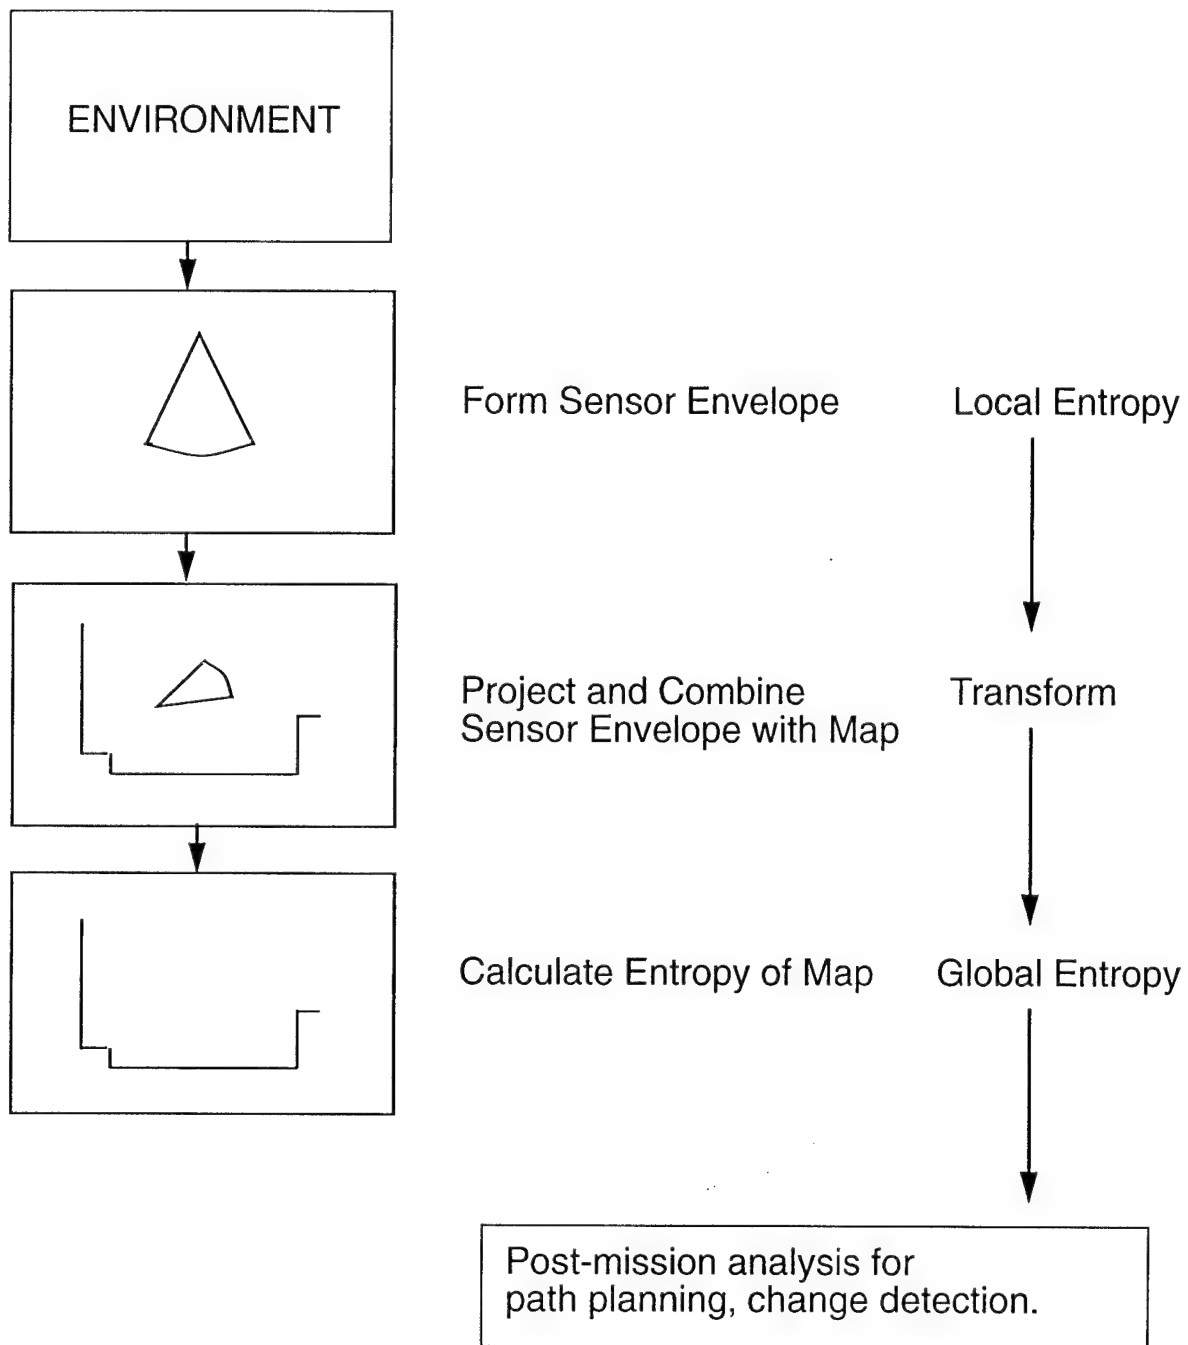


Figure 6.1 Post Mission Analysis

A stable convergent map should have the following characteristic: the majority of the pixels that have been mapped should have converged to stable consistent values. We note however that the convergence of a map is measured on a global scale and that global convergence does not imply convergence for individual pixels. Thus we choose to low-pass filter the convergent map, which allows us to look at neighbourhoods of pixels rather than the somewhat noisier individual pixels. Figure 6.2 shows the entropy calculated for the pseudo-2D map constructed in section 5.3 where we have calculated the entropy for four pixel by four pixel averaged neighbourhoods. We analyze the characteristics of this image in the next section.

6.2 Entropy for Path Planning and Change Detection

If we look at the convergent map (Fig 6.2) after our mapping exercises it shows a number of interesting features. The most obvious feature is the terrain that has not been explored. Part of this terrain is marked A. Terrain may be unexplored due to sensor limitations or due to the fact that our mapping exercises have not progressed to that particular region.

Within the regions that have been mapped, we see that for our stable convergent map the vast majority of the pixels have converged to stable values (areas labelled B1 and B2). We note that the stable regions include regions that may be occupied — such as around the pilings corresponding to the MBL dock marked B1; or empty — such as the area labelled B2.

Other interesting features are also labelled C, D, and E. The area marked C contains evidence for occupancy but has not converged to that value. If we compare this region with the corresponding region in figs 5.13 through 5.18 we see that region C was not considered occupied for the first 1600 pings. In fact an examination of our data shows that just before we stopped our mapping

efforts we received a few hits from this region which point towards occupancy. On the one hand this reinforces the fact that global convergence does not guarantee convergence at a local level, but on the other a temporal analysis of this region shows the possibilities of using our methodology for change detection. We believe that the change in occupancy occurred due to the presence of a school of fish in the area and it is encouraging to see that we can identify such temporally dynamic events.

The area marked D is interesting in so far that it contains unexplored regions but the regions around it have been explored and have converged to consistent values. As can be seen in the Figs 5.15-5.18, which plot out the occupied pixels, this region is one that has been occluded by the presence of an object in the field of view of the sensor.

Lastly, the area marked E indicates one region that is occupied but has not yet converged to a stable value. This is primarily due to the fact that our operations were confined close to the MBL dock and the spreading effects of the beam (at the long ranges associated with the pilings in E) tend to smear out these features, thereby preventing us from localizing them consistently.

Given this analysis we are now in a much better position to plan subsequent operations in the area. We might simply opt to map the unexplored regions. More interestingly, if our interests were confined to the areas already under consideration, we realize that interesting regions include the areas marked C, D, and E. Region C is interesting from the standpoint of detecting change in our environment; region D is interesting because we have very little information about this region, while region E contains features that we realize have not been adequately mapped out.




Figure 6.2 Entropy for Convergent Map

CHAPTER 7

Conclusions

In this thesis we examined the task of high-frequency acoustic mapping from an underwater vehicle. High-frequency acoustic mapping is fundamentally a difficult task because it involves the interaction of quantities of very high dynamic range.

7.1 Contributions of this thesis

We demonstrated in this thesis that:

- A stochastic sensor modelling framework may be defined based on the characteristics of a particular sensor such that stochastic backprojection provides us with consistent mapping performance.
- We can model the dominant uncertainties in our mapping paradigm by the use of the concept of entropy. Our formulation of entropy for stochastic modelling possesses properties that confirm our intuition about acoustic mapping.
- In using an entropic framework it is useful, at times, to consider the entropies of occupancy and emptiness separately, as our years of intuition in working with acoustic sensors is based upon our experiences in two-dimensional mapping.
- Under typical operating conditions we are limited in the performance of our mapping algorithms primarily by the navigational and attitudinal uncertainties of the vehicle from which the acoustic sensor is being utilized and by the sampling constraints of the sensor and its ability to

interrogate the environment. A formulation is derived for looking at optimizing vehicle parameters with respect to the information garnered during the entire mission.

- A fully three-dimensional solution for stochastic world modelling is presented that respects the computational constraints of typical operations.

- The three-dimensional solution for stochastic world modelling was used within an entropic framework to demonstrate the utility of this work for acoustic mapping in a challenging real-world environment. Our formulation was seen to be stable, converging, and excellent for typical underwater mapping tasks, such as focussing on objects of interest.

- We also showed the utility of our entropic formulation for post-mission analysis with a view to doing path-planning for subsequent missions.

7.2 Directions for Future Work

This thesis answers a number of questions about the nature of high frequency acoustic mapping but important issues in the field still need to be addressed. Very little work has been done so far in the area of spatiotemporal mapping. Geologists are very interested in characterizing the rate of change of events at hydrothermal sites. Mapping work carried out in the Guayamas Basin [93] showed the utility of such work. To accomplish such work with an AUV will requires us to characterize the uncertainties in our mapping over long periods of time. We need also to be able to use our high-frequency sonars in an autonomous mode. This will require repackaging and redesigning some of the data handling and storage capabilities so as to deploy these high-frequency sensors from an AUV, but technically such redesign is feasible today. At a higher level, given the dynamic nature of the acoustic mapping we need to understand the mechanisms for adjusting the gains optimally on these sensors. This work is crucial towards achieving the goal of using high-frequency

acoustic sensors on AUVs for long-term deployments at a site.

The entropic framework can easily be merged with existing modelling programs for understanding the operating envelopes of underwater vehicles in different environmental conditions. Significant work is being carried out in this area [26][113]. It has focussed on the graphical, human interface, and computational issues of displaying existing data. These simulations can be easily extended and merged with the work presented in this thesis for the purposes of probabilistic analysis of target detectivity, mission planning, and post-mission analysis.

As pointed out earlier, we require a better understanding and characterization of AUV energy consumption. Plans are already being made to carry out this task [99]. In general, we should begin to get a much better handle on some of these parameters as AUVs transition from an experimental to operational status. This work should allow us to use the basic sensor modelling framework laid out in this thesis to optimize higher level mission goals.

The good performance of our algorithm in focussing on objects of interest opens the door to work in terrain-relative navigation [41] as well as for servoing on objects of interest. A significant limitation of underwater vehicles is the need for external navigation systems. The promise of feature-based navigation opens up new possibilities for AUVs. Given the greater ranges that are foreseen with these vehicles, the way we make measurements in the ocean could be revolutionized.

Work also should be carried out on incorporating global stochastic modelling to Kalman filter based vehicle-centered representations of the environment [32] to exploit the strengths of each method. We note that an entropic formulation is not limited solely to use with global modelling. It could just as easily be used with the local feature-tracking Kalman filter algorithms.

We also need to examine the performance of our stochastic modelling methodology in dead-

reckoning based navigation systems, where the errors in navigation are a function of the total distance travelled.

Finally, we point out that an entropic formulation is not specifically limited to acoustic mapping. A thorough understanding of the physics behind a particular sensing modality should allow one to model that sensing modality in terms of entropy and provide the end user with a way of quantifying the degradations in sensor performance for that particular sensing modality.

APPENDIX A: Derivation of the Range Equation for AUVs

This derivation follows that given by Bradley[18]. We begin by assuming that the drag force (F_d) on the vehicle is

$$F_d = (1/2)\rho C_d S u^2 \quad (\text{A.1})$$

where ρ is the density of water, C_d is the drag coefficient, S is the surface area for the body, and u is the speed at which the vehicle is travelling.

The energy E to travel a distance R is just

$$E = F_d R \quad (\text{A.2})$$

For a battery pack of mass m_b and motor propeller efficiency η , the energy available for propulsion is

$$E = \eta B_0 m_b \quad (\text{A.3})$$

where B_0 is the energy density of the batteries in J / kg.

With a floatation material having a specific gravity σ , the mass that can be carried is

$$m_b = V\rho(1 - \sigma) \quad (\text{A.4})$$

where V is the volume of the floatation material. We can express the volume V and the surface area S of the AUV as functions of its length L .

$$V = V_0 L^3 \quad (\text{A.5})$$

$$S = S_0 L^2 \quad (\text{A.6})$$

where V_0 and S_0 are non-dimensional parameters which depend upon the shape of the vehicle.

Substituting values from (A1), (A4), (A5), and (A6) into (A2) and (A3) we get

$$\eta B_0 V_0 L^3 \rho (1 - \sigma) = (1/2) \rho C_d S_0 L^2 u^2 R \quad (\text{A.7})$$

which reduces to

$$R = \left(\frac{2 (1 - \sigma) V_0}{C_d S_0} \right) \frac{B_0 L}{u^2} \quad (\text{A.8})$$

Equation (A.8) gives the range for an AUV when we consider the propulsion load alone. If we consider the hotel load P_H for sensing then Eqn (A.2) is modified to

$$E = F_d R + P_H (R/u) \quad (\text{A.9})$$

Following the previous analysis the range then becomes

$$R = \left(\frac{2 (1 - \sigma) V_0}{C_d S_0 + \left(\frac{2 P_H}{\rho L^2 u^3} \right)} \right) \frac{B_0 L}{u^2} \quad (\text{A.10})$$

We can maximise R in (A.10) with respect to u , to see that there is only one maxima and that that maxima occurs at a velocity

$$u = \left(\frac{P_H}{2 C_d S_0 \rho L^2} \right)^{1/3} \quad (\text{A.11})$$

APPENDIX B: KALMAN FILTERING FOR UNDERWATER VEHICLES

B.1 Kalman filter formulation

Mapping requires some way of measuring the uncertainty in vehicle attitude and navigation. We use a Kalman filter [54] for this purpose. The Kalman filter has been used extensively for merging measurements from disparate sources. A full discussion of Kalman filtering for underwater vehicles is beyond the scope of this thesis. In this appendix we briefly lay out the basic formulation of a particular realization of a Kalman filter that was used for estimating navigation and attitude for the results presented in the next chapter. Further details of the actual Kalman filter implementation may be found in a forthcoming report [108] and a related article [72].

Our state vector includes states corresponding to three-dimensional vehicle accelerations, velocities, and positions as well as the pitch, roll, heading, pitch rate, roll rate, and yaw (rate of change of heading).

The Kalman filter provides an unbiased, minimum variance, consistent estimate of a quantity \underline{x} , based on a set of measurements \underline{z} , where the relationship between these quantities can be represented by

$$\underline{z} = H\underline{x} + \underline{v} \quad (\text{B.1})$$

where \underline{v} represents random noise.

The notation we follow is outlined below.

$\underline{x}(t|t)$ is the estimate of the state vector using all data up to and including the present time step.

$\underline{x}(t|t-1)$ is the predicted estimate of the state vector at the time of the next filter update step

using all data up to but not including the present time step.

$\underline{P}(t|t)$ is the estimate of the state vector using all data up to and including the present time step.

$\underline{P}(t|t-1)$ is the predicted estimate of the state vector at the time of the next filter update step using all the data up to but not including the present time step.

We derive error covariance matrices based on sensor characteristics as determined by experiments in the field.

The governing equations for the system model are

$$\dot{\underline{x}} = A\underline{x} + G\underline{w} \quad (\text{B.2})$$

$$\underline{z} = C\underline{x} + \underline{v} \quad (\text{B.3})$$

where \underline{w} is a random forcing function and the matrices A and G relate the terms in the equation.

The update step in the filter propagation consists of calculating the innovation term I , the covariance matrix of innovation \underline{v} , the Kalman gain K , the new state vector estimate $\underline{x}(t|t)$ and finally the new error covariance $\underline{P}(t|t)$ according to the equations below.

$$I = \underline{z} - (C \times \underline{x}(t|t-1)) \quad (\text{B.4})$$

$$\underline{v} = C \times P(t|t-1) \times C' + R \quad (\text{B.5})$$

$$K = P(t|t-1) \times C' \times \text{inv}(\underline{v}) \quad (\text{B.6})$$

$$\underline{x}(t|t) = \underline{x}(t|t-1) + K \times I \quad (\text{B.7})$$

$$P(t|t) = (\text{Ident} - K \times C) \times P(t|t-1) \times (\text{Ident} - K \times C)' + K \times R \times K' \quad (\text{B.8})$$

where $\text{inv}()$ refers to the inverse of a matrix, Ident is an identity matrix with the appropriate dimensions, and R is the covariance matrix of measurement noise.

The predicted values of the state vector $\underline{x}(t+1|t)$ and the error covariance matrix

$\underline{P}(t+1|t)$ are now calculated. As the Kalman filter is implemented as a discrete filter we need also to convert the continuous matrices A and G to their discrete counterparts ϕ and Γ . After these conversions are made we move to the prediction step as outlined below.

$$\underline{x}(t+1|t) = \phi \times \underline{x}(t|t) + \Gamma \times \underline{w} \quad (\text{B.9})$$

$$P(t+1|t) = \phi \times P(t|t) \times \phi' + \Gamma \times Q \times \Gamma \quad (\text{B.10})$$

where Q is the error covariance matrix for the white noise \underline{w} .

B.2 Navigation and Attitude Sensors

In the following sub-sections we briefly discuss typical sensors for navigation and attitude. We also discuss methods for minimizing and quantifying the errors associated with these sensors.

B.2.1 Navigation

For the datasets that we discuss in this thesis, the two primary navigation systems used were the SHARPS [124] and a low-frequency long baseline (LBL) system. The SHARPS has been shown to give repeatable measurements between fixes to 2 cm. The accuracy achievable with such a system allows us to use multiple transducers on the vehicle to solve problems with obtaining a reliable heading reference.

The ultimate limitation in using such a system is the dynamics of the vehicle in the time between pings. For a highly maneuverable ROV like JASON this could result in much poorer performance during turns. In acoustic mapping, typically we do not merge the data collected during turns into our map for these reasons.




Figure B.1 Variance in Kalman Filtered Navigation

The other navigation system used for mapping exercises with ABE, is a long-baseline system, which has operating characteristics such that the repeatability is within a few meters.

The performance of these systems is usually not uniform across the entire net. The errors are influenced by the shadowing effects of bottom bathymetry, by the directional behavior of the transponder's beam patterns, by sound velocity profile effects, and the geometry of the net.

In the z dimension, temperature compensated pressure transducers [87] provide nominal accuracies of 0.02% over the entire range. This translates to accuracies of tens of cm in deep water and much better performance in shallower water.

The problem of data dropout between fixes can be attacked by using an Acoustic Doppler Current

Profiler (ADCP) [89]. This low-power instrument was originally designed for studying bottom currents but can also be used in a bottom-tracking mode to obtain a very reliable velocity reference. Rowe [92] reports accuracies of 1% of total distance travelled in using such a sensor in the field. In our own experience, we can easily achieve accuracies of 5% of distance travelled as shown in Fig B.1. The disparity between our results is primarily because of errors in our heading reference.

B.2.2 Attitude

For pitch and roll on JASON we use a Watson dual-axis inclinometer [119]. This uses an angular rate sensor combined with a precision vertical reference. The angular rate sensor consists of piezoelectric bender elements mounted to a rigid base in a tuning-fork configuration. The two drive forks are driven in opposite directions. The momentum stored in the vibrating elements during rotation causes an out-of-plane bending force, which is demodulated to accurately depict the rotation rate. This rate is integrated to obtain a position reference, which is combined with the vertical reference to obtain accurate attitude measurements.

Figure B.2 plots out the pitch and roll values from the sensors while the vehicle was at rest on its cradle. The roll value is seen to be quite stable although there is seen to be some change in the pitch value due probably to the manner in which the vehicle was resting in its cradle.




Figure B.2 Pitch / Roll for Vehicle at Rest

A cheap, power efficient, robust heading reference is hard to achieve. The primary sources of error in a flux-gate magnetic compass [69][69] are caused by dynamic magnetic fields in the operating environment. These can at best provide a low-frequency source of heading information. They may be seriously compromised in a shallow-water environment or geological regions of magnetic anomaly.

The other possible source of heading reference is some form of free-spinning gyro [110][110]. North-seeking gyros are both too expensive and too bulky to be used on AUVs. The free spinning gyros maintain excellent short term stability but tend to drift at rates of $5\text{-}10^\circ$ per hour.

B.3 Heading reference

A cheap, power efficient, robust heading reference is hard to achieve. The primary sources of error in a flux-gate magnetic compass [69] are caused by dynamic magnetic fields in the operating environment. These can at best provide a low-frequency source of heading information. They may be seriously compromised in a shallow-water environment or geological regions of magnetic anomaly.

The other possible source of heading reference is some form of free-spinning gyro [110]. North-seeking gyros are both too expensive and too bulky to be used on AUVs. The free spinning gyros maintain excellent short term stability but tend to drift at rates of 5-10° per hour.

A decent heading reference is crucial for reducing the errors associated with integrating acceleration and velocity references for dead reckoning. It is also crucial for high frequency acoustic mapping. A magnetic flux-gate compass mounted on an ROV is very susceptible to the effects of stray magnetic fields, such as those associated with the vehicle's thrusters and metallic frame. While some of these stray effects may be calibrated out, some of the magnetic effects, such as the time-varying magnetic fields induced by the thrusters, severely limit the bandwidth of typical magnetic compasses.

Obviously, one must try to minimize the magnetic interference from thrusters by shielding them appropriately as well as by placing the magnetic compass at a location where the effects are minimized. But the truth unfortunately remains that a magnetic compass cannot provide the high-frequency (that is in the 2-5 Hz range) data that is required for high frequency acoustic mapping.

Another form of heading reference is the gyroscopic angular-rate sensors. These use a vibrating quartz tuning fork to sense angular rate by acting as a Coriolis sensor. These can form the basis of

providing a precise angular measurement but tend to drift over time. The gyro angular rate sensors that are affordable from an ROV or an AUV operational standpoint have typical drifts on the order of $10^\circ / \text{hr}$.

The solution to obtaining a cheap heading reference is to use some form of complementary filter to merge the low-frequency accuracy of the fluxgate compass with the high-frequency component of a gyro.

Another option exists for operations within a reliable navigation net. If the vehicle is performing a survey pattern, a series of navigation fixes may be used to derive a heading for the vehicle. This may then be used as a low-frequency source to be combined with a gyro to obtain a good heading reference. One must be wary in using such a course of action, however, because flyers may significantly affect the performance of the algorithm. However, in some areas, such as at geological sites with high magnetic anomalies and in shallow water in the presence of docks, some form of navigation-derived heading may be necessitated.

For the dataset in this thesis we ignored the magnetic compass completely, choosing instead to use the dual very high precision SHARPS navigation transponders on the vehicle to derive a heading. This heading was used to compute a gyro bias, which was subtracted from the instantaneous gyro reading to provide an accurate high-frequency heading reference.

Figure B.3 plots the gyro bias versus time. The plot shows that from an unknown state the gyro bias quickly converges to the true value then follows the drift of the gyro quite well.




Figure B.3 Gyro Bias vs Time

APPENDIX C: Geometric modelling for 3-D range sensing

Geometric modelling may *not* be suitable for autonomous mapping operations because of the noisy nature of high-frequency acoustic mapping. However, it may be used quite successfully for ROV operations where operators can monitor data quality.

In using a geometrical modelling approach one should check for data consistency by examining redundant coverage of the same area from different look angles. The x,y,z values are typically processed with a gridding package such as GMT [121][100]. The processing steps may include a block-median, a block-mean, and binning before a surface is fitted to the gridded data. Smoothness constraints on the surface tend to low-pass filter the dataset. The low-pass nature of these images may be taken to be representative of our navigational uncertainties.

The procedure for geometric modelling is straightforward. Careful measurements must be made of the positions and orientations of each of the sensors. One must also be careful about calibrating the attitude sensors to avoid inconsistencies that may arise from attitude sensor biases. The process becomes one of transforming the data from each sensor from the sensor coordinate frame to the vehicle coordinate frame and from there to a frame referenced to the world. The results of such an exercise are shown in Fig C.1 where we show the plan view for our three dimensionally, geometrically gridded, dataset. Matlab source code to geometrically grid the pencil-beam sonar data is included at the end of this appendix.




Figure C.1 Geometrically Gridded Data

```

% plot_meso_xyz.m - geometrically grids the x,y,z data
%   hanu           - 12-24-94
%
load ../data/meso.ascii      % contains the julian time / meso_ang / meso_range
load meso_state_all         % contains state estimates for vehicle

DTOR = pi / 180;

MES_Xo = 0;                % Sonar sensor offsets
MES_Yo = 0;
MES_Zo = 0;
PAR_Zo = 0;                % Depth (Paroscientific) sensor offsets

% NOTE: Assumed Sensor angular offsets about x and y were
% measured to be less than 0.1 degrees

non_zero_indices = find(meso(:,3)<65.0 & meso(:,3)>2.0);
meso_nz = zeros(length(non_zero_indices),3);

for i = 1:length(non_zero_indices),
    meso_nz(i,:) = meso(non_zero_indices(i,:),:);
end;

start_sample = 100;
end_sample = 5990;
figure(1);
clg;
% axis([-20 70 -30 60]);
hold on;

time_index = 1;
% skip to the relevant time
for j = start_sample:end_sample,
    while(new_ALL(time_index,1) < meso_nz(j,1))
        time_index = time_index + 1;
    end;

    % Now get interpolated values for state vars
    dt_state = state(time_index,1) - state(time_index - 1,1);
    dt = state(time_index - 1,1) - meso_nz(i,1);

    dx = state(time_index,4) + state(time_index - 1,4);
    nav_x = (dx * dt_state/dt) + state(time_index - 1,4);
    dy = state(time_index,5) + state(time_index - 1,5);
    nav_y = (dy * dt_state/dt) + state(time_index - 1,5);
    dz = state(time_index,6) + state(time_index - 1,6);
    nav_z = (dz * dt_state/dt) + state(time_index - 1,6);

    dpitch = state(time_index,10) + state(time_index - 1,10);
    att_pitch = ((dpitch * dt_state/dt) + state(time_index - 1,10)) ;
    droll = state(time_index,11) + state(time_index - 1,11);

```

```

att_roll = ((droll * dt_state/dt) + state(time_index - 1,11)) ;
dhdg = state(time_index,12) + state(time_index - 1,12);
att_hdg = ((dhdg * dt_state/dt) + state(time_index - 1,12)) ;
% setup rotation matrices to convert from sensor to vehicle frame
att_hdg = 2*pi - att_hdg;
pcos = cos(att_pitch);
psin = sin(att_pitch);
rcos = cos(att_roll);
rsin = sin(att_roll);
hcos = cos(att_hdg);
hsin = sin(att_hdg);

M(1,1) = rcos*hcos;
M(2,1) = hsin*rcos;
M(3,1) = -rsin;
M(1,2) = psin*rsin*hcos - pcos*hsin;
M(2,2) = psin*rsin*hsin + pcos*hcos; M(3,2) = psin*rcos;
M(1,3) = pcos*rsin*hcos + psin*hsin;
M(2,3) = pcos*rsin*hsin - psin*hcos;
M(3,3) = pcos*rcos;

% Transform sensor to world
nav_z = nav_z + PAR_Zo;
mes_x = nav_x + M(1,1)*MES_Xo + M(1,2)*MES_Yo + M(1,3)*MES_Zo;
mes_y = nav_y + M(2,1)*MES_Xo + M(2,2)*MES_Yo + M(2,3)*MES_Zo;
mes_z = nav_z - M(3,1)*MES_Xo + M(3,2)*MES_Yo + M(3,3)*MES_Zo;

dmang = meso_nz(j,2)*DTOR;

% Calculate components for forward looking sonar
D = [sin(dmang) cos(dmang) 0];

% Calculate direction cosines of sensor cone
p_fcos(1) = M(1,1)*D(1) + M(1,2)*D(2) + M(1,3)*D(3);
p_fcos(2) = M(2,1)*D(1) + M(2,2)*D(2) + M(2,3)*D(3);
p_fcos(3) = M(3,1)*D(1) + M(3,2)*D(2) + M(3,3)*D(3);
p_fcos(3) = -p_fcos(3);

% Calculate grid points
grid_x(j) = meso_nz(j,3)*p_fcos(1) + mes_x;
grid_y(j) = meso_nz(j,3)*p_fcos(2) + mes_y;
grid_z(j) = meso_nz(j,3)*p_fcos(3) + mes_z;
end;

% plot plan view of gridded points and vehicle tracks
plot(grid_x,grid_y,'.');
plot(state(start_sample:end_sample,4),state(start_sample:end_sample,5),'r')

```

REFERENCES

- [1] Arkin, R., Motor-schema-based mobile robot navigation, *International Journal of Robotics Research*, vol. 8, no. 4, 1989.
- [2] Ayache, N., Faugeras, O., Building, registering and fusing noisy visual maps, *International Journal of Robotics Research*, vol. 7, no. 6, pp. 45-65, 1988.
- [3] Ayache, N., Faugeras, O., Maintaining representations of the environment of a mobile robot, *IEEE Transactions on Robotics and Automation*, vol. 5, no. 6, pp. 804-819, 1989.
- [4] Bailey, H.S., The Voyage of the Challenger, *Scientific American*, pp. 88-94, May 1953.
- [5] Ballard, R.D., *The Discovery of the Titanic*, Warner Books Inc., New York, 1987.
- [6] Ballard, R. D., D. R. Yoerger, W. K. Stewart, and A. Bowen, ARGO/JASON: A Remotely Operated Survey and Sampling System for Full-Ocean Depth, in *Proceedings of the IEEE Oceans 91 Conference*, pp. 71-75, Sept. 1991.
- [7] Barshan, B., and R. Kuc. Differentiating sonar reflections from corners and planes by employing an intelligent sensor, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(6), pp. 560-569, June 1990
- [8] Beckerman, M. and E. Oblow. Treatment of systematic errors in the processing of wide-angle sonar sensor data for robotic navigation, *IEEE Trans. Robotics and Automation*, vol. 6, no. 2, pp. 136-145, 1990.
- [9] Berkson, J. M., and J. E. Matthews, Statistical characterization of seafloor roughness, *IEEE Journal of Oceanic Engineering*, Vol. 9, pp. 48-51, 1984.
- [10] Bernoulli, J., *Ars Conjectandi*, Basel, 1713. Reprinted in 1968 by Culture et Civilisation, 115 Avenue Gabriel Lebon, Brussels.
- [11] Bellingham, J.G., Design and Utilization of Small, High Performance Autonomous Underwater Vehicles, *Commercialization of Autonomous Underwater Vehicles*, January 25-26, 1994, Cambridge, Massachusetts.
- [12] Bellingham, J., Goudey, C.A., Consi, T.R., Bales, J.W., Atwood, D.K., Leonard, J.J., Chrysostomidis, C., A Second Generation Survey AUV, *Proceedings of the 1994 Symposium on Autonomous Underwater Vehicle Technology*, Cambridge, Massachusetts, July, 1994.
- [13] Blackington, J. G., Bathymetric resolution, precision and accuracy considerations for swath bathymetry mapping sonar systems, *Proceedings of the IEEE Oceans 91 Conference*, pp. 550-557, 1992.
- [14] Blackington, J.G., Hussong, D.M., Kosalas, J.G., First Results from a Combination Side-scan Sonar and a Seafloor Mapping System (SeaMARC II), *Proceedings of the 15th Annual Offshore Technology Conference*, vol. 1, pp. 307-314, Houston, Texas, May 1983.

- [15] Blidberg,D.R., Allmendinger,E.E., Sideris,N., The Development of an Unmanned, Self-controlled, Free-Swimming Vehicle, *Proceedings of the 10th Offshore Technology Conference, Houston, TX, 1978.*
- [16] Blidberg,D.R., Time-ordered Architecture for Knowledge Based Guidance of an Unmanned Untethered Submersible, *Proceedings of the MTS/IEEE Oceans '84 Conference, 1984.*
- [17] Blidberg,D.R., Guidance Control Architecture for the EAVE Vehicle, *IEEE Journal of Oceanic Engineering*, vol. 11, no. 4, pp. 449-461, October, 1986.
- [18] Bradley, A., Low power navigation and control for long range autonomous underwater vehicles.
- [19] Bradley, A., Duester, A.R., Liberatore, S., A Fast Hydrographic Profiling System, *Proceedings of the MTS/IEEE Oceans '91 Conference, 1991.*
- [20] Bolle, R. M., and D. B. Cooper, On optimally combining pieces of information, with application to estimating 3-D complex-object position from range data, *IEEE Transactions on Pattern Analysis and Machine Intelligence*. vol.8, no. 5, pp. 619-638, 1986
- [21] Borenstein, J., and Y. Koren, Obstacle avoidance with ultrasonic sensors, *IEEE J. Robotics and Automation*, vol. RA-4, pp. 213-218, April 1988.
- [22] Borenstein, J., and Y. Koren, Real-time obstacle avoidance for fast mobile robots in cluttered environments, *IEEE International Conference on Robotics and Automation, 1990.*
- [23] Brooks, R. A., A robust layered control system for a mobile robot, *IEEE J. Robotics and Automation*, vol. RA-2, no.1, pp14-23, , March 1986.
- [24] Brown, M. K., Feature extraction techniques for recognizing solid objects with an ultrasonic range sensor, *IEEE J. Robotics and Automation*, vol. RA-1, no. 4, pp 191-205, December 1985.
- [25] Brown, M. K., The extraction of curved surface features with generic range sensors, *Int. J. Robotics Research*, vol. 5, no. 1, pp 3-18, 1986.
- [26] Brutzman, D., Kanayama, Y., Zyda, M., NPS AUV Integrated Simulator, *Proceedings of IEEE AUV 1992 Conference, June 1992.*
- [27] Cady, W., Q-Flex(R) Accelerometer Construction and Principle of Operation, Technical Note TN-103, Sundstrand Data Control, Inc., Redmond, WA, 1984.
- [28] Catipovic, J.A., Design and performance analysis of a digital acoustic telemetry system, *Sc.D. thesis, MIT, Cambridge, MA, 1987.*
- [29] Clay, C. S., H. Medwin, and W. M. Wright, Specular scattered sound and the probability density function of a rough surface, *Journal of the Acoustical Society of America*, Vol. 53, pp. 1677-1682, 1973.

- [30] Cobra, D., Estimation and Correction of Geometric Distortions in Side-Scan Sonar Design . *Ph.D. Thesis*, MIT-WHOI Joint Program in Oceanography, June 1990.
- [31] Cox, I. J., J. B. Kruskal, and D. A. Wallach, Predicting and estimating the performance of a subpixel registration algorithm, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1990.
- [32] Cox, I.J., Leonard, J.J., Modelling a dynamic environment using a Bayesian multiple hypothesis approach, *Artificial Intelligence*, vol. 66, pg. 311-344, 1994.
- [33] Crowley, J.L., World modeling and position estimation for a mobile robot using ultrasonic ranging, *Proceedings of the IEEE International Conference on Robotics and Automation*, pp. 674-681, 1989.
- [34] D'Appolito, J., The Evaluation of Kalman Filter Designs for Multisensor Integrated Navigation Systems, The Analytic Sciences Corporation, *Technical Report AFAL-TR-70-271*, 1970.
- [35] Davis, D.L., Pilskaln, C.H., Measurements with Underwater Video: Camera Field Width Calibration and Structured Lighting, *Marine Technology Society Journal*, vol. 26, no. 4, pp. 13-19, Winter 1992-1993
- [36] de Moustier, C., Beyond bathymetry: Mapping acoustic backscattering from the deep seafloor with Sea Beam, *Journal of the Acoustical Society of America*, Vol. 79, pp. 316-331, 1986
- [37] de Moustier, C., and H. Matsumoto, Seafloor Acoustic Remote Sensing with Multibeam Echo-Sounders and Bathymetric Sidescan Sonar System, *Marine Geophysical Research.*, Vol. 15, no. 1, pp. 27-42, 1993.
- [38] Dempster, A.P., Upper and lower probabilities induced by a multi-valued mapping, *Annals of Mathematical Statistics*, vol 38, pp. 325-339, 1967.
- [39] Denbigh, P. N., , A bathymetric sidescan sonar, in *Proceedings of the Ultrasonics International 79 Conference*, pp. 321-326, 1979.
- [40] Denbigh, P. N., Swath bathymetry: Principles of operation and an analysis of errors, *IEEE Journal of Oceanic Engineering*, Vol. 14, pp. 289-298, 1989.
- [41] DiMassa, D., Concurrent Mapping and Navigation in an Unengineered Underwater Environment, *Ph.D. Thesis Proposal, Dept of Ocean Engineering, MIT, Cambridge, Massachussets*, 1994.
- [42] Drumheller, M. Mobile robot localization using sonar, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-9, pp. 325-332, March, 1987.
- [43] Dunsiger, A. D., N. A. Cochrane, and W. J. Vetter, Seabed characterization from broadband acoustic echosounding with scattering models, *IEEE Journal of Oceanic Engineering*, Vol. 6, pp. 94-106, 1981.

- [44] Dyer, I., *Class Notes for 13.851: Principles of Underwater Sound*, MIT, Cambridge, MA, 1990.
- [45] Eckart, C., The scattering of sound from the sea surface, *Journal of the Acoustical Society of America*, Vol. 25, pp. 566-570, 1953.
- [46] Egeskov, P., Bjerrum, A., Pascoal, A., Silvestre, C., Aage, C., Smitt, L.W., Design, Construction and Hydrodynamic Testing of the MARIUS AUV, *Proceedings of the 1994 Symposium on Autonomous Underwater Vehicle Technology*, Cambridge, Massachusetts, July, 1994.
- [47] Elfes, A. Sonar-based real world mapping and navigation, *IEEE Journal of Robotics and Automation*, vol. RA-3, no. 3, pp. 249-265, June 1987.
- [48] Ewing, M., Vine, A.C., Worzel, J.L., Recent Developments in Submarine Geophysics, *Geological Society of America Bulletin*, vol. 57, pp 909-934, 1946.
- [49] Farr, H. K., Multibeam bathymetric sonar: Sea Beam and Hydrochart, *Marine Geodesy*, vol. 4, pp. 77-93, 1980.
- [50] Flynn, A. M., Combining sonar and infrared sensors for mobile robot navigation, *International Journal of Robotics Research*, vol. 7, no. 6, 1988.
- [51] Fox, C. G., and D. E. Hayes, Quantitative methods for analyzing the roughness of the seafloor, *Rev. Geophys.*, Vol. 23, pp. 1-48, 1985.
- [52] Grimson, W.E.L., *From Images to Surfaces: A Computational Study of the Human Early vision System*, MIT Press, Cambridge, Massachusetts, 1981.
- [53] Grimson, W. E. and T. Lozano-Perez, Model-based recognition and localization from sparse range or tactile data, *International Journal of Robotics Research*, vol. 3, no.3, pp3-35, 1984.
- [54] Gelb, A., ed., *Applied Optimal Estimation*, MIT Press, Cambridge, MA, 1974.
- [55] Harmon, S. Y., The ground surveillance robot (GSR): An autonomous vehicle designed to transit unknown terrain, *IEEE Journal of Robotics and Automation*, vol. RA-3, no. 3, pp 266-279, June 1987.
- [56] Horn, B.K.P., Brooks, M.J., eds. *Shape from Shading*, The MIT Press, Cambridge, MA, 1989.
- [57] Hoyt, J.K., The Flying Fish, an Untethered Oceanographic Sensor Platform with Acoustic Homing Capability, *Ph.D. Thesis, MIT-WHOI Joint Program*, May 1986. WHOI-86-18.
- [58] Jackson, D. R., A. M. Baird, J. J. Crisp, and P. A. G. Thomson, High-frequency bottom backscatter measurements in shallow water, *Journal of the Acoustical Society of America*, vol. 80, pp. 1188-1199, 1986.

- [59] Jackson, D. R., D. P. Winebrenner, and A. Ishimaru, Application of composite roughness model to high-frequency bottom backscattering, *Journal of the Acoustical Society of America.*, vol. 79, pp. 1410-1422, 1986.
- [60] Jackson, D. R., and K. B. Briggs, High-frequency bottom scattering: Roughness versus sediment volume scattering, *Journal of the Acoustical Society of America.*, vol. 92, pp. 962-977, 1992.
- [61] Jaffe, J., Computer Modelling and the Design of Optimal Underwater Imaging Systems, *IEEE J. of Oceanic Engineering*, vol. OE-15, no. 2, pp. 101-111, April 1990.
- [62] Jalving, B., The NDRE-AUV Flight Control System, *IEEE Journal of Ocean Engineering*, vol.19, no.4, pp. 497-501, October 1994.
- [63] Jarvis, R., A perspective on range finding techniques for computer vision, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 5, no. 2, pp 122-139, March 1983.
- [64] Klein, M., Side Scan Sonar, *Undersea Technology*, vol. 8, pp. 24-26, 1967.
- [65] Klepsvik, J.O, Torsen, H.O., Thoresen, K. Laser Imaging Technology for Subsea Inspection: Principles and Applications, *IRM/ROV 1990 Conference*, November 1990.
- [66] Kojima, J., Ito, Y., Asakawa, K., Shirasaki, Y., Kato, N., Development of Autonomous Underwater Vehicle for Inspection of Underwater Cables, *Proceedings of Underwater Intervention '93*, pp 349-356, 1994.
- [67] Kuc, R., Siegel, M.W., Physically based simulation model for acoustic sensor robot navigation, *IEEE Transactions on Pattern Analysis and Machine Intelligence*. vol. PAMI-9, no. 6, pp. 766-778, November, 1987.
- [68] Kusters, J. G., Analysis and application of an underwater optical ranging system, *Engineer Thesis*, Massachusetts Institute of Technology/Woods Hole Oceanographic Institution, 1992.
- [69] KVH Industries Inc., *C-100 Compass Engine Technical Manual*, Rev. D, Middletown, RI, 1992.
- [70] Lambert, J.H., *Neues Organon*. Reprinted in 1965 as the first two volumes of Lambert's Philosophische Schriften by Georg Olms Verlagsbuchhandlung, Hildesheim.
- [71] Langer, D., Rosenblatt, J.K., Hebert, M., A Behavior-Based System for Off-Road Navigation, *IEEE Transactions on Robotics and Automation*, vol. RA-10, no. 6, December 1994.
- [72] Leader, D.E., Kalman Filter Estimation of Underwater Vehicle Position and Attitude Using a Doppler Velocity Aided Inertial Motion Unit, *Engineer's Thesis*, MIT-WHOI Joint Program, September 1994.
- [73] Leonard, J., Directed Sonar Sensing for Mobile Robot Navigation, *Ph.D. Thesis*, Oxford University, 1990.

- [74] Malik, S., Quantitative backscatter estimation using a bathymetric sidescan sonar, *Master Thesis*, U. Virginia, 1991.
- [75] Malinverno, A., Segmentation of topographic profiles of the seafloor based on a self-affine model, *IEEE Journal of Oceanic Engineering*, vol. 14, pp. 348-359, 1989.
- [76] Malmquist, K. A., Modelling a 300-kHz bathymetric sonar system, *Master Thesis*, Massachusetts Institute of Technology/Woods Hole Oceanographic Institution, 1992.
- [77] Marks, R.L., Rock, S.M., Lee, M.J., Real-time Video Mosaicking of the Ocean Floor, *Proceedings of the 1994 IEEE Symposium on Autonomous Underwater Vehicle Technology*, Cambridge, MA, July, 1994.
- [78] Masnadi-Shirazi, M. A., C. de Moustier, P. Cervenka, and S. H. Zisk, Differential phase estimation with the SeaMARC II bathymetric sidescan sonar system, *IEEE Journal of Oceanic Engineering*, vol. OE-17, pp. 239-251, 1992.
- [79] Mitchell, N. C., and M. L. Somers, Quantitative backscatter measurements with a long-range sidescan sonar, *IEEE Journal of Oceanic Engineering*, vol. 14, pp. 368-374, 1989.
- [80] Moran, B., Underwater Shape Reconstruction in Two Dimensions, *Ph.D. Thesis*, Department of Ocean Engineering, MIT, Cambridge, MA, 1994.
- [81] Morgan, M., *Dynamic Positioning of Offshore Vessels*, The Petroleum Publishing Company, Tulsa, 1978.
- [82] Moravec, H., Personal Communication, 1993.
- [83] Moravec, H. Sensor fusion in certainty grids for mobile robots, *Sensor devices and systems for Robotics*, pp. 253-276, Springer-Verlag, 1989. NATO ASI Series.
- [84] Mourad, P. D., and D. R. Jackson, High-frequency sonar equation models for bottom backscattering and forward loss, in *Proceedings of the IEEE Oceans 89*, pp. 1168-1175, 1989.
- [85] Negahdaripour, S., Yu, C.H., Shokrallahi, A., Recovering Shape and Motion from Undersea Images, *IEEE Journal of Oceanic Engineering*, vol. 15, no. 3, pp. 189-198, July, 1990.
- [86] Papoulis, A., *Probability, Random Variables, and Stochastic Processes*. Mc-Graw Hill, Inc., New York, NY. 1984.
- [87] Paroscientific Inc., *Standard Operating Procedure, Doc. No. 8099*, Redmond, WA, 1987.
- [88] POLAROID *Ultrasonic Ranging System Handbook*, Application Notes / Technical Papers, 1984.
- [89] RD Instruments, *Direct-Reading Broadband Acoustic Doppler Current profiler Technical Manual*, Rev. 4, San Diego, 1993.

- [90] Reed, T. B., D. Chu, P. R. Shaw, D. Smith, and T. Stanton, A comparison of methods for quantitative analysis of sonar imagery and bathymetry, *1991 Fall Research Symposium on Bottom/subbottom Acoustic Reverberation Special Research Program*, Sept. 24-25, 1991.
- [91] Reed, T. B., and D. Hussong, Digital image processing techniques for enhancement and classification of SeaMARC II sidescan sonar imagery, *Journal of Geophysical Research*, vol. 94, pp. 7469-7490, 1989
- [92] .Rowe, F., Brumley, B., Integrated Navigation System Development, *Phase I Final Report*, RD Instruments, San Diego, 1992.
- [93] Robigou, V., Stewart, K., Ballard, R., Hydrothermal Vent Sites in the Guayamas Basin revisited by an ROV: High Precision Bathymetric, Geological and Biological Mapping. Invited Presentation, *EOS Transactions of the American Geophysical Union*, vol. 74, pg. 573, 1993.
- [94] Rusby, S., A long range side-scan sonar for use in the deep sea (G.L.O.R.I.A. Project), *International Hydrographic Review*, vol. 47, no. 2, pp. 25-39, 1970.
- [95] Russell-Cargill, W.G.A, ed. *Recent Developments in Side Scan Sonar Techniques*, ABC Press, Capetown, South Africa, 1982.
- [96] Saridis, V., On the theory of intelligent controls, *Proceedings of the SPIE Conference on Intelligent Robots and Computer Vision*, vol 848, pp. 488-495, Cambridge, MA. November 1987.
- [97] Shafer, G., *A Mathematical Theory of Evidence*, Princeton University Press, 1976.
- [98] Shannon, C.E., *The Mathematical Theory of Communication*, University of Illinois Press, Urbana, IL. 1948.
- [99] Singh, H., Bradley, A., Stewart, W.K., Yoerger, D., Sonar Mapping with the Autonomous Benthic Explorer (ABE), *In preparation*.
- [100] Smith, W.H.F., Wessel, P., Gridding with continuous curvature splines in tension, *Geophysics*, vol. 55, pp. 293-305, 1990.
- [101] Somers, M.L., Carson, R.M., Revie, J.A., Edge, R.H., Barrow, B.J., Andrews, A.G., GLORIA II - An Improved Long Range Side-scan Sonar, *Proceedings of IEE-IERE Subconference on Ocean Instruments and Communications*, Oceanology International, B.P.S. Publishers Ltd., London, pp. 16-24, 1978.
- [102] Stanton, T. K., Sonar estimates of seafloor microroughness, *Journal of the Acoustical Society of America*, Vol. 75, 809-818, 1984.
- [103] Stanton, T., Chu, D., Martin, L., Forthcoming publication.
- [104] Stewart, W.K., High Resolution Optical and Acoustic Remote Sensing for Underwater Exploration, *Oceanus*, vol. 34, no. 1, pp 10-22, Spring 1991.

- [105] Stewart, W.K. Multisensor modelling underwater with uncertain information, *Tech. Rep. AI-TR 1143*, MIT AI Lab, Cambridge, MA, 1988.
- [106] Stewart, W.K., D. Chu, X. Tang, A Preliminary Study of Shallow-Water Sonar Issues: Signal Motion Loss and Reverberation Noise, *WHOI Tech Report WHOI-93-41*.
- [107] Stewart, W.K., Chu, D., Malik, S., Lerner, S., Singh, H., Quantitative Seafloor Characterization using a Bathymetric Sidescan Sonar, *IEEE Journal of Oceanic Engineering*, vol. 19, no. 4, pp. 599-610, October, 1994.
- [108] Stewart, W.K., Tang, X., Singh, H., DiMassa, D., Kalman Filtering for the JASON ROV, *WHOI Technical Report*, Forthcoming.
- [109] Stewart, W. K., M. Marra, and M. Jiang, A hierarchical approach to seafloor classification using neural networks, in *Proceedings of the IEEE Oceans 92 Conference*, Oct. 1992.
- [110] Systron Donner Inertial Division, *Gyrochip(tm) Theory of Operation*, Concord, CA, 1993.
- [111] Tucker, M.J., Stubbs, A.R., A narrow-beam echo-ranger for fishery and geological investigations, *British Journal of Applied Physics*, vol. 12, pp. 103-110, 1961.
- [112] Tuohy, S.T., Geophysical Map Representation, Abstraction and Interrogation for Autonomous Underwater Vehicle Navigation, *Ph.D. Thesis, Department of Ocean Engineering, MIT*, Cambridge, Massachusetts, 1993.
- [113] Tuohy, S., A Simulation Model for AUV Navigation, *Proceedings of the 1994 IEEE Symposium on Autonomous Underwater Vehicle Technology*, Cambridge, MA, July, 1994.
- [114] Tusting, R.F., Davis, D.L., Laser Systems and Structured Illumination for Quantitative Undersea Imaging, *Marine Technology Society Journal*, vol 26, no. 4, Winter 1992-1993.
- [115] Ura, T., Free Swimming Vehicle PTEROA for Deep Sea Survey, *Proceedings of the ROV 1989*, pp 263-268, 1989.
- [116] Urick, R. J., *Principles of Underwater Sound*, McGraw-Hill, New York, 1975.
- [117] Urick, R. J., and R. M. Hoover, Backscatter of sound from the sea surface: its measurement, causes and application to the prediction of reverberation levels, *Journal of the Acoustical Society of America*, Vol. 28, pp. 1038, 1956.
- [118] von Alt, C., Allen, B., Austin, T., Stokey, R., Remote Environmental Measuring Units, *Proceedings of the 1994 Symposium on Autonomous Underwater Vehicle Technology*, Cambridge, Massachusetts, July, 1994.
- [119] Watson Industries Inc, *Watson Inclinator with Angular Rate Owner's Manual*, Eau Claire, WI, 1990.
- [120] Walsh, D., Undersea Satellites: The Commercialization of AUVs, *Marine Technology Society Journal*, vol. 27, no. 4, pp. 64-73, 1994.

- [121] Wessel, P., Smith, W.H.F., Free Software helps map and display data, EOS Transactions of the American Geophysical Union, vol. 72, pp. 441, 445-446, 1991.
- [122] Williams, A.J., Personal Communication.
- [123] Winther, S., *SPOTSCAN 2D System Description, Offshore Trials and Test Report*, Seatex AS, Trondheim, Norway, December 1990.
- [124] Yoerger, D., Mindell, D., Precise Navigation and Control of an ROV at 2200 metres depth, *Proceedings of the MTS ROV/Intervention '92 Conference*, 1992.
- [125] Yoerger, D., Robotic Undersea Technology, *Oceanus*, vol. 34, no. 1, pp 32-37, Spring 1991.
- [126] Yoerger, D., A. Bradley, Dynamic testing of the autonomous benthic explorer.
- [127] Yoerger, D., A. Bradley, B. Walden, The autonomous benthic explorer (ABE): A deep ocean AUV for scientific seafloor survey.
- [128] Zadeh, L.A., A Simple View of the Dempster-Shafer Theory of Evidence and its Implication for the Rule of Combination, *AI Magazine*, vol. 7, no. 2, pp. 85-90, Summer 1986.
- [129] Zorpette, G., Unmanned Vehicles Make a Splash, *IEEE Spectrum*, vol. 31, no. 8, pp. 38-44, August 1994.

DOCUMENT LIBRARY

Distribution List for Technical Report Exchange - May 1995

University of California, San Diego
SIO Library 0175C
9500 Gilman Drive
La Jolla, CA 92093-0175

Hancock Library of Biology & Oceanography
Alan Hancock Laboratory
University of Southern California
University Park
Los Angeles, CA 90089-0371

Gifts & Exchanges
Library
Bedford Institute of Oceanography
P.O. Box 1006
Dartmouth, NS, B2Y 4A2, CANADA

Commander
International Ice Patrol
1082 Shennecossett Road
Groton, CT 06340-6095

NOAA/EDIS Miami Library Center
4301 Rickenbacker Causeway
Miami, FL 33149

Research Library
U.S. Army Corps of Engineers
Waterways Experiment Station
3909 Halls Ferry Road
Vicksburg, MS 39180-6199

Institute of Geophysics
University of Hawaii
Library Room 252
2525 Correa Road
Honolulu, HI 96822

Marine Resources Information Center
Building E38-320
MIT
Cambridge, MA 02139

Library
Lamont-Doherty Geological Observatory
Columbia University
Palisades, NY 10964

Library
Serials Department
Oregon State University
Corvallis, OR 97331

Pell Marine Science Library
University of Rhode Island
Narragansett Bay Campus
Narragansett, RI 02882

Working Collection
Texas A&M University
Dept. of Oceanography
College Station, TX 77843

Fisheries-Oceanography Library
151 Oceanography Teaching Bldg.
University of Washington
Seattle, WA 98195

Library
R.S.M.A.S.
University of Miami
4600 Rickenbacker Causeway
Miami, FL 33149

Maury Oceanographic Library
Naval Oceanographic Office
Building 1003 South
1002 Balch Blvd.
Stennis Space Center, MS, 39522-5001

Library
Institute of Ocean Sciences
P.O. Box 6000
Sidney, B.C. V8L 4B2
CANADA

Library
Institute of Oceanographic Sciences
Deacon Laboratory
Wormley, Godalming
Surrey GU8 5UB
UNITED KINGDOM

The Librarian
CSIRO Marine Laboratories
G.P.O. Box 1538
Hobart, Tasmania
AUSTRALIA 7001

Library
Proudman Oceanographic Laboratory
Bidston Observatory
Birkenhead
Merseyside L43 7 RA
UNITED KINGDOM

IFREMER
Centre de Brest
Service Documentation - Publications
BP 70 29280 PLOUZANE
FRANCE

REPORT DOCUMENTATION PAGE	1. REPORT NO. MIT/WHOI 95-14	2	3. Recipient's Accession No.
4. Title and Subtitle An Entropic Framework for AUV Sensor Modelling			5. Report Date June 1995
			6
7. Author(s) Hanumant Singh			8. Performing Organization Rept. No.
9. Performing Organization Name and Address MIT/WHOI Joint Program in Oceanography/Applied Ocean Science & Engineering			10. Project/Task/Work Unit No. MIT/WHOI 95-14
			11. Contract(C) or Grant(G) No. (C) N00014-90-C-0098 (G)
12. Sponsoring Organization Name and Address Office of Naval Technology			13. Type of Report & Period Covered Ph.D. Thesis
			14.
15. Supplementary Notes This thesis should be cited as: Hanumant Singh, 1995. An Entropic Framework for AUV Sensor Modelling. Ph.D. Thesis. MIT/WHOI, 95-14.			
16. Abstract (Limit: 200 words) <p>This thesis examines the general task of active sensing by defining a measure of efficiency for sensing in a particular environment. We focus on fine-scale acoustic mapping from an autonomous underwater vehicle (AUV). The constraints on imaging underwater – vehicle power, vehicle hydrodynamics, computational and telemetry requirements, and typical navigational and attitudinal uncertainties along with the underlying physics of the acoustic sensing modality – are considered in defining an entropic measure of sensor efficiency.</p> <p>675-kHz pencil-beam sonar data acquired using the JASON remotely operated vehicle in a challenging shallow water environment and 200-kHz echo-sounder data acquired using the ABE AUV are used to demonstrate the utility of the entropic framework.</p> <p>We show the utility of an entropic framework for the following: (i) Optimizing the speed of the AUV for maximizing the information gathered with a particular sensor. (ii) the rate of convergence and the stability of our mapping efforts in the face of typical uncertainties in navigation and attitude; (iii) as a methodology for actual sensor deployment and use on a real vehicle; and (iv) in tasks such as post-mission analysis for applications such as change detection and path planning for subsequent missions.</p>			
17. Document Analysis			
a. Descriptors			
b. Identifiers/Open-Ended Terms			
c. COSATI Field/Group			
18. Availability Statement Approved for publication; distribution unlimited.	19. Security Class (This Report) UNCLASSIFIED	21. No. of Pages 136	
	20. Security Class (This Page)	22. Price	